Quality Technology Solutions, Inc.

10 Steps to Implement a Disaster Recovery Plan

A Review of Strategies and Options for Implementing a Disaster Recovery Strategy

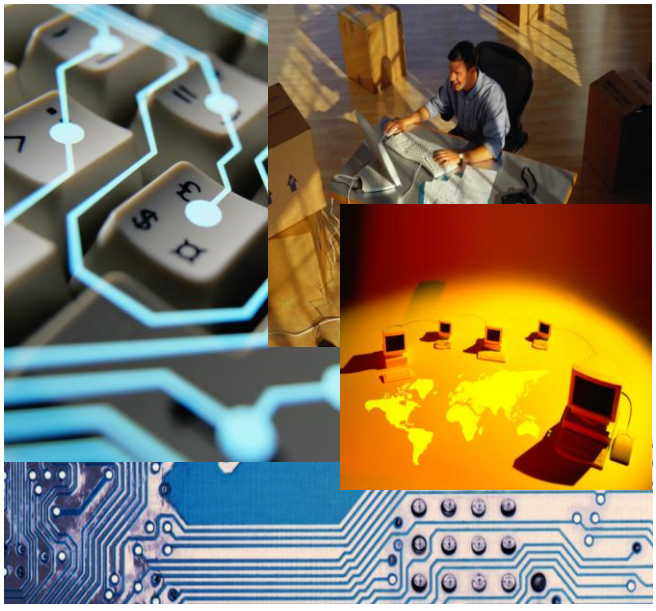

Authored by Neil A. Rosenberg, a Certified Information Systems Security Professional ("CISSP") and frequent speaker and author on information security and disaster recovery topics, this paper will discuss a suggested methodology for developing and implementing a Disaster Recovery Plan, with particular focus on medium-sized organizations. With increased business reliance on technology, disaster recovery has become a critical matter for business survival. In this white paper, we will provide a strategy for helping you to protect your business assets.

Providing Worry-Free Network Solutions!

Quality Technology Solutions, Inc.
1639 Route 10, Suite 103
Parsippany, NJ 07054

(973)984-7600
www.QTSnet.com
info@QTSnet.com

Table of Contents

Why Do We Care?

Disaster Recovery Planning is a relevant topic for businesses of all sizes and types. It all comes down to mitigating risk for your business.

Various studies done over the years offer different numbers, but generally the result is the same – somewhere between 70 and 90 percent of all businesses that suffer from a disaster are out of business within 2 years of the disaster.

The reasons for this are fairly intuitive. A disaster scenario casts doubt on the viability of the business to survive the incident, among the company's customers, supply chain, employees, shareholders and others.

Businesses with a well-defined plan have the ability to maintain that confidence even in trying times, and restore normal operations quickly. Those that do not, cannot. And the impact is often lethal.

Introduction

This white paper is intended to give its readers a framework with which to review their business networks against a baseline of ten fundamental and important elements of implementing a Disaster Recovery Plan. Some businesses may have Disaster Recovery Plans in place, and in that case this white paper provides a useful checkpoint, and hopefully validation.

It seems like many prescriptive lists boil down to "ten points," and after some careful review I am presenting what I have seen to be the ten key steps to building and implementing a Disaster Recovery Plan for a medium business (and to some degree both a small business and enterprise) network. I would say that they represent ten "best practices" that should generally be applicable across the spectrum.

Once you have your plan, with sequenced activities and a realistic timeline, you are ready to begin. And I am confident that by following this process, you will see a genuine and substantive improvement in your organization's disaster preparedness, and a subsequent reduction in business risk.

Since the majority of medium business networks are based on Microsoft technology, we'll pay particular attention to Microsoft products here, but these points apply across the board.

Disaster Recovery versus Business Continuity Planning

Although these terms are often used interchangeably, Disaster Recovery Planning is distinct from Business Continuity Planning. The key distinction between these different, but related, activities would be:

- Business Continuity Planning is about defining the assets, threats and scenarios that can adversely impact your organization, and making decisions about how or to what degree to mitigate these risks. In other words, Business Continuity Planning is in most cases about ***preventing the Disaster Scenario from happening.***
- Disaster Recovery Planning is about defining consistent, pre-planned actions that will react to various Disaster Scenarios. In other words, Disaster Recovery Planning is about ***reacting to the Disaster Scenario after it has happened***.

Business Continuity Planning centers on the business as a whole, and business processes. Disaster Recovery Planning often starts in IT, not because that is the only place of focus but because it is the most obvious

Business Continuity Planning & Business Impact Analysis

Business Continuity Planning is a comprehensive and often complex discipline that delves deeply into the business as a whole. In fact, IT really is in most cases a support function to the business, and many businesses have separate departmental systems for specific business functions.

Business Continuity Planning often starts with a ***Business Impact Analysis***, which is similar to the Asset/Threat discussion noted in this White Paper, but tends to go much more deeply into the economic impact on the business. Usually this is a hard economic analysis, and drives deeper into the workflow of different business units. These studies often take weeks or months, and lead to a very comprehensive understanding of business impact, and the DR Plan.

(since systems are centralized) and the easiest (because business units can abdicate responsibility and say Disaster Recovery is "an IT thing." This attitude often puts IT organizations in a position of having a business-wide responsibility, without business-wide backing or financial resources.

Although these are different activities, they are clearly related and Disaster Recovery Planning builds upon Business Continuity Planning when properly done.

QTS' QuikRecover™ methodology, based on our focus on the medium business market space, integrates elements of these two processes into one engagement. Our process is not a comprehensive Business Continuity Planning engagement, nor do we do a comprehensive, economic Business Impact Analysis – for a more comprehensive (and expensive) engagement, consult with your accounting firm or a specialist in this area.

1. Define Key Assets, Threats and Scenarios

A Disaster Recovery or Business Continuity effort should start with identification of key assets, and definition of the business impact of loss of each asset. This is a critical step – you need to know what you're protecting and what its value to the business is, to define how it should be protected. We do this routinely in our QuikSecure Security Policy engagement and our QuikRecover Disaster Recovery Plan Workshop engagement – essentially, you can pull your business team together in a conference room and brainstorm on all the key assets of the business – electronic, paper, physical, etc. – and identify the impact of potential loss. This exercise allows you to determine, based on the impact of loss, the appropriate protection required for each asset.

Examples of Assets to be considered in your planning include:

- Accounting/ERP System and related components (order entry, inventory, warehouse management, etc.);
- Email system and archives (or, at a more basic level, the ability to send and receive email);
- Files and Documents on your Local Area Network;
- Product Designs and Specifications (written/electronic);
- Business Strategy Documents (written/electronic);
- Facilities and Fixed Assets;
- Paper Files and Documents;
- Cash and Valuables on Premise;
- Employee Knowledge and Staff Expertise (human resources);
- Product Inventory and Raw Materials.

Management Buy-In

Disaster Recovery Planning and Business Continuity Planning are business issues, and decisions that need to be made by business management.

This discussion is really about business risk, not technology. Each asset has a business impact, and therefore a value which is worth protecting with a certain level of investment. When this process is IT-driven, the appropriate business buy-in is seldom accomplished, and cost is not weighed against risk.

You can deal with risk in three different ways:

- ***Accept*** the risk and do nothing;
- ***Assign*** the risk (insurance);
- ***Avoid*** the risk by taking mitigating actions, such as developing a Disaster Recovery Plan.

This is clearly a business (not just IT) discussion.

Another key part of this process is evaluating threats against your business, and each asset. An organization with a single office in New York City, for example, does not need to be very concerned about the threat of earthquakes, but does have more concern around terrorism than an organization headquartered in Wyoming. Organizations that are geographically diverse have a wider range of threats based on this.

Potential Threats to be considered and mitigated against could include:

Natural Threats

- Fire
- Floods and Flash Floods
- Hurricane/Storm
- Earthquake
- Blizzard/Severe Winter Storm

Man-Made Threats

- Terrorism
- Explosion
- Bio-Hazard
- Epidemic/Pandemic
- Theft/Vandalism
- Work Stoppage
- Riot
- Power/HVAC Failure
- Communications Failure
- Hardware Failure
- Software Failure
- Security Incident

Scenarios can then be defined, based on different threats at each location. Those scenarios (for example, a site outage where the facility is still intact) should have pre-defined actions providing guidance on whether a disaster should be declared and whether the Disaster Recovery site should be activated, as well as define the recovery window (see below).

2. Determine The Recovery Window

Once you have defined your assets, you need to determine how long you can go without access to this resource. This is called the “Recovery Window” for each asset. Clearly the recovery plan is going to be very different (and less expensive) if your IT assets can be unavailable for 3-5 days following disaster, versus a mandate of being up within 3-5 hours.

Application Chains

QTS defines the concept of "Application Chains" as a critical element of Disaster Recovery Planning. An Application Chain consists of all of the elements of IT infrastructure associated with delivery of an IT service or business asset.

For example, the asset "Email" may actually consist of many components, any one of which could bring down access to email as a company asset:

- Email Server
- Email Gateway
- Firewall
- Switches
- Anti-Spam Service
- Internet Access
- Internet Router
- DNS Server

Based on an understanding of the IT components that support each asset, one can determine and eliminate Single Points of Failure for High Availability as well as Disaster Recovery Planning.

Since this difference potentially has a significant cost impact, it is important that this decision – asset by asset – be tied to the business value of the asset. "Knee jerk" reactions such as "everything needs to be available in hours" are not only unrealistic, but don't allow the IT team to set and manage against priorities. Working under pressure in a disaster situation, it is important for the people executing against the plan to have clear direction and priorities. Some of your systems may have a 1 hour (or less threshold), while others may be fine if they are operational the next day, and this allows the IT team to focus on the most important systems first. This requires consensus and input from your management team – hence the value of the exercise in the prior step, when done properly.

3. Define Recovery Solutions

The third step builds upon the first two, and here we define the appropriate approach and solutions based on the assets and the Recovery Window. Solutions can include recovering from tape backup or disk backup, or data replication to an offsite location with "hot" failover. Determining the appropriate type and level of protection ties directly to the business value of the asset, and how long you can work without it.

For example:

- An e-commerce web site may need to be operational at all times since it is customer-facing, and this dictates co-location and possibly data replication.
- Email may need to be available within hours – but depending on the requirements, the solution could vary between replication of the whole email system versus "turning on" a secondary web-based email site.
- An accounting system might be fine if the Finance staff can be up and running in 24-48 hours, which possibly maps to a tape backup/restore solution.

Each Asset, based on its defined value and Recovery Window, can then have an appropriate Disaster Recovery solution identified, with a commensurate budget that maps back to business value and impact.

For example, Microsoft offers its Exchange Hosted Continuity Service, which may be more cost effective for smaller businesses than replicating the entire Exchange message store to a DR site. Or, even if it is not, such a solution might allow IT resources to be focused on other systems.

Salvage Team & Operations

Although recovery gets most of the focus in Disaster Recovery planning, Salvage operations are an important element. Someone needs to assess the damage to the main facilities and what resources can be salvaged. In most cases, the damage to a site is only partial, and judgment is required here.

Equally important, at some point the business will want to resume normal operations out of its main offices, whether that is days, weeks or months after the disaster. Salvage planning and the ability to get the site ready for restored operations is very important, and requires dedicated resources.

It is important to accurately assess the resources at the impacted facilities, to recover whatever can be salvaged, and to properly plan re-establishment of operations at these facilities.

4. Draft A Disaster Recovery Plan

Based on these preliminaries, it is now time for the fourth step - developing a draft Disaster Recovery Plan. This plan will be defined by the assets and how they will be protected, but will also address key process and communication-related elements. Furthermore, the process for assessing damage to the existing site, as well as mitigating/minimizing damage, needs to be considered.

There are many logistical considerations. For example, how will employees get to the Disaster Recovery site, and what if employees cannot get to the site? In the case of severe human loss at the primary site or travel/transportation problems, this can be a real problem. Some careful thought needs to go into defining the true emergency plan for recovery and staffing of key functions.

An important element of the Disaster Recovery Plan is establishing an Emergency Operations Center – a location where employees and partners responsible for executing on the Disaster Recovery Plan can converge to in order to work together and have access to resources and decision-making authority. This location should be pre-established and have resources including office supplies, telecommunications, food and water, communications lines and other needed tools. This is often, but not always, at the hot site utilized for the data center.

The most important rule of Disaster Recovery Planning is “People First.” In other words, the safety of your employees and personnel is always a more important consideration than saving assets or business recovery operations. It is important that your team understand this, and that this be a consistent element of your planning and message.

5. Establish A Communications Plan and Assign Roles

Our fifth step is actually a key subset and component of the Disaster Recovery Plan: identifying the communication plan and assigning roles and responsibilities to member of your Disaster Recovery team.

Some key issues to be considered include:

- Who is responsible for “declaring” a disaster, and what is the communication chain? Does IT simply assume it should fail over systems, or is this a business decision – not all disasters are obvious, either in scope or duration. For example, if there is a building-wide

power failure, how long does the business wait before switching over to backup systems?
- What is the communication plan to employees?
- What is the communication plan to customers? To suppliers and partners? To the general public?
- Who is authorized to communicate with shareholders? With governmental agencies including municipal (fire, police) and regulatory?
- Which Disaster Recovery team members are responsible for recovery operations (getting the Disaster Recovery site up and running), versus salvage operations (evaluating the current site to define what can be recovered)?

From a communications standpoint it is important to have a well-organized, accurate and up-to-date list of contacts for each function or role, including emergency contact home and cell phone numbers, email addresses (remember, corporate email may be down) and a chain of communications so instructions can be distributed hierarchically – it is important to be able to move quickly, and efficient communication is the key to this.

All of this requires some level of planned thought, and management buy-in. Equally important, who is responsible for each element, and who are the backup personnel? Are they trained, and are there "out of band" communication channels to initiate key communications and activities? In many cases, going through the exercise shows that organizations have an unrealistic dependency on key people, who need to execute many important tasks concurrently – and often impossibly. The planning process helps define those situations so responsibilities can be allocated in a realistic manner.

Employee training is absolutely critical. Once people know what their roles are (or in fact that they even have roles), it is important that they know how to execute on their tasks, who to escalate issues to or seek instruction from, and where to find the most current version of the plan. If people are not trained, and clear on their responsibilities, the actual disaster scenario will go very badly.

Many organizations do not have the resources in-house to deal with all of the different roles and elements of executing the Disaster Recovery Plan, which means that outside organizations (such as systems integrators, telecommunications providers and hardware vendors) may play key roles in executing on the Disaster Recovery Plan. When that is the case, it is critical that these vendors understand their role in the plan, and that service level agreements (or at least clear understandings of responsibility) are in place

Telecomm Considerations

Voice and Data Communications is a critical element of any Disaster Recovery Plan. Sometimes, in organizations where IT does not have responsibility for phone system, this gets overlooked.

For voice systems, it is important to have a failover plan to be able to switch over services to an alternate location or system in the event of loss of the primary site, or failure of the phone system at the primary site. Fault tolerance/ High Availability need to be planned for the phone system as much as data systems.

Fax also need to be considered, whether using POTS lines or integrated into the phone system.

Data comm can often be failed over using BGP for Internet circuits, or meshing services like Frame Relay. ISDN dial backup is giving way to meshed, switched fabrics.

Co-Location and Branch Office Sites

Although traditional data center backup operations like Sungard, HP (formerly Comdisco) and IBM are doing quite well with the increased focus on disaster recovery, many businesses are looking to newer web co-location providers, or their own branch offices, for hosting of DR sites.

Part of this is a matter of control. Some businesses found after 9/11 that if they were not the first to declare a disaster and reserve space, they were down at the bottom of the priority list and sometimes did not get service. In other cases, these alternate approaches can save money.

Many web hosting organizations have seen opportunity in the DR space, and already have the needed infrastructure, presenting a cost-effective alternative to traditional DR vendors.

so those vendors are properly prepared and resourced to assist you with execution. Don't just assume that they will be, or can be.

6. Disaster Recovery Site Planning

Assuming the above 5 steps have been put in place, the next step is to implement the systems or capabilities required to deliver the plan. In most cases, this will involve definition of some type of disaster recovery site to deal with situations where the data center or main facility are unavailable. There are three primary types of sites:

- A "Hot Site" is a site that has live communication links, working systems and real data (typically with real-time data replication), ready for an immediate or near-immediate failover of operations;
- A "Warm Site" typically has live communication links and some amount of hardware, but typically requires installation of software and/or restoration of data from tape or another media format – typically in a span of hours or a day before the site is operational;
- A "Cold Site" typically is a facility where your staff can go to when a disaster is declared, and which may have external communications, but which does not typically have pre-existing hardware, software or data.

Obviously, the cost associated with a Hot Site is vastly different from that of a Warm or Cold Site – hence the importance of the initial steps defining importance, business value/impact to thus determine the suitable investment level to properly protect the business.

There are several ways to implement data replication in a Hot Site scenario, ranging from timed copying of files or database backups (least automated/expensive) to using replication software such as Double-Take Software's Double-Take, to SAN-to-SAN hardware-based data replication (the most expensive and complex).

In any replication scenario, it is VERY important to account for the bandwidth requirements for the data being copied from site to site. Data replication can consume significant bandwidth, and therefore it is important to plan this and ensure that disaster recovery replication does not consume precious Internet access bandwidth or WAN links and compete with existing applications. WAN acceleration, via products such as Citrix WANScaler, can provide a valuable option to get the most out of existing links without needing to purchase extra bandwidth or new links (but this varies situation by situation). Timing of replication can also be an issue – for example, it might be acceptable to throttle communications during the work day, and

Co-Location and Branch Office Sites (continued)

When planning a Disaster Recovery Hot or Warm site, it is important to consider a wide range of factors:

- Different power grid, geographically disparate to main office;
- Adequate Power and HVAC capacity;
- High-speed telecomm links with dedicated IP addresses (redundant if possible);
- Adequate space for servers, PCs, users, Emergency Operations Center;
- Proximity to key employees and to transportation (public transit, highway access);
- Access to transportation services;
- Availability of human resources.

If all of these can be properly accounted for at a branch office or Co-Location site, then this can be a viable approach.

have replication "catch up" after 6pm if the Recovery Window allows for loss of a few hours' data.

Some organizations (at the higher end of cost and infrastructure) implement what is sometimes called a "stretch cluster" where the two nodes of a server cluster are split across two sites. This is very hard to implement largely due to the maturity of the technologies and the communications challenges – even if sufficient bandwidth is available for the communications, latency increases as servers are spread further apart. Normally, this will only work with Fiber Optic connections, which have inherent distance limits. This technology will no doubt become more robust in future years.

7. Accessing Data and Applications

The seventh step is defining the appropriate mechanism for access to your data and applications. In the old days (not so long ago, really), this typically meant setting up rows of PCs and phones at work tables and desks, for staff to work from the Disaster Recovery site. This creates some very significant logistical issues – since you want the Disaster Recovery site to be geographically disparate from the main location(s), this makes it very difficult for employees to get to the Disaster Recovery site. In many cases, you will need to arrange transportation for the employees to the site. This gets more complex when you look at the likelihood of also needing to transport families (due to child care responsibilities), which raises the alternate challenge of spouses who may have conflicting responsibilities. Being realistic, it is harder and harder to relocate employees, even temporarily, to Disaster Recovery sites.

With the prevalence of universal connectivity, the more modern approach is to have some computers and work space at the Disaster Recovery site, but the majority of users connecting via VPN from home computers, laptops or remote sites. Some applications are web-server based and therefore lend themselves to this approach, but most applications even today are 32-bit Windows-based, and for these applications technologies like Citrix MetaFrame Presentation Server and Windows Terminal Services allow these "thick" applications to be run over low-bandwidth links from thin client systems. In effect, this approach extends the Disaster Recovery site worldwide if needed, and significantly improves the capability of businesses to integrate more employees into the Disaster Recovery Plan as needed.

Failover of inbound communications needs to be carefully planned – will all inbound traffic be re-routed to the Disaster Recovery site, including email, remote access, web access, etc.? How will this be initiated and who will handle it? Or will there be a separate access method to the Disaster

Tape Considerations

Tape Backup is the most common form of disaster recovery. Most businesses assume that they can recover their data by simply backing up to tape. Some are sadly mistaken.

It is absolutely critical to rotate tapes offsite. Consider your exposure when planning offsite tape rotation. If you don't rotate your tapes offsite daily, you are subject to multiple days of data loss in the event of a disaster.

When keeping some of your tapes onsite, store them vertically to avoid the pull of gravity warping your tapes. Tapes should be stored in a fire-rated (not just fire-proof) safe that can withstand the heat that goes with fire (most fire-proof safes prevent fire from getting in, but not necessarily the heat that can damage tapes).

Also, perform test restores on at least a monthly basis, to make sure you can actually do a tape restore.

Recovery site? Border Gateway Protocol ("BGP") failover solutions like AT&T's MARO offerings can be used to automate failover, or MX records can be set with different weightings for email, or tools like Citrix' NetScaler appliances can be used to re-route traffic dynamically. Careful thought needs to go into planning the right solution for your business.

8. Document The Disaster Recovery Plan, In Detail

The eighth step is an absolutely critical one – document the plan. Although you developed a draft Disaster Recovery plan in step 4 above, in this step it is important to develop a more detailed plan for each system on exactly what to do to implement failover to the Disaster Recovery site system. This is actually something QTS refers to as System Recovery Plans, which are subset components of a Disaster Recovery Plan.

Your Disaster Recovery Plan, and component System Recovery Plans, need to be very specific and detailed – in the middle of a crisis, you don't want to leave much open to interpretation, and you may not be sure of the identity or skill level of the people implementing the plan. As we noted above, the people implementing the plan might even not be your employees, or might not be the first-line employees responsible for that task.

An important element of System Recovery Plans is documentation of how to implement failback to the original systems once they become available again. This can include messy tasks like data synchronization and re-load of servers at the main facility. The complexity of these tasks can make it a tough call as to when to declare a disaster – for example, in the case of a power outage when you don't know if you will be back up in an hour or a day, is it worth the effort of moving operations, then moving them back? This is why the whole process requires some planning and thought/discussion before the disaster occurs.

9. Test The Disaster Recovery Plan

Now, it is time to Test The Plan! Only through a real test will the staff be familiar with what to do in the event of a disaster, but equally important, through this process you can identify the gaps, inconsistencies and errors in the plan. It is much better to identify problem areas in a test than in a real disaster. Again, you don't want to leave things open to interpretation in the middle of a pressurized, stressful situation.

Disaster Recovery Plan Elements

Your Disaster Recovery Plan needs to be tailored to the unique needs and circumstances of your business. However, some recommended elements of the DR Plan and supporting documents would include:

- Asset Matrix
- Threats Matrix
- Scenarios
- Disaster Mitigation Procedures (by Scenario)
- Backup Procedures
- Application Chains
- Disaster Recovery Plan Activation Procedure
- Communications Chain and Plans
- System Recovery Plans for each IT Component supporting noted Assets
- Business Recovery Plans for each Business Unit (including departmental IT systems)
- Salvage Plan
- Staff Training & Awareness
- DR Plan Maintenance Procedures

Prior to an actual, live test of the plan, you should perform a "checklist test." With the plan being documented in detail, the Disaster Recovery team should sit down as a group and review the plan documents step by step, to identify weaknesses and missing steps. This way, the plan will be as thorough as possible before the actual first test.

This first test is absolutely critical, and should be performed with a complete failover and failback of your systems. Follow the plan carefully – ideally, the execution is done by personnel who did not write the documents, so assumed steps aren't done. This presents an opportunity to further refine the plan, identifying additional missing steps and weaknesses. It also allows for weaknesses in employee training to be identified, and remedied. Obviously, this needs to be done in a window of time where actual risk to your business, and disruption of the business, is minimized. A review meeting should be held promptly afterward to debrief on what went well and what did not, and to drive needed refinements and changes into the plan and into your business processes.

Testing of the Disaster Recovery Plan should be re-done annually, as systems and the business change. Some would say that tests should be a surprise to your user community, but that applies more to the human/staffing elements of facility evacuation plans rather than IT plans.

It is also important that the Disaster Recovery Plan be kept current and "fresh" as your systems change. A good way to do this is to integrate the Disaster Recovery Plan into your Change/Control process, or at least set quarterly checkpoints to review the plan and determine what updates need to be made.

10. Refine and Re-Test the Disaster Recovery Plan

Finally, the 10th Step is to refine the plan and its documentation, and conduct a re-test based on this revised and refined version. The second test should be much smoother than the first, and should put you in a position where you are ready to execute against the plan in a real disaster recovery situation. Again, the re-test should be an annual activity in your IT business plan.

Finally, ensure that you store your plan offsite at the Emergency Operations Center, along with copies of all software media and license keys, copies of system documentation and other useful documents. It is also important for key VAR or systems integration partners to have current copies of your plans. This way, everyone can be on the same page and ready to react as a team in the case that it really becomes necessary to do so.

Highlights:

- Frequent Computer Industry Speaker
- Author of numerous information security articles for various publications
- 20+ years of computer industry experience
- Numerous technical certifications including CISSP
- Extensive consulting and project management experience
- Attorney

About the Author

Neil A. Rosenberg is the President and CEO of Quality Technology Solutions, Inc., a leading New Jersey network integration and security services firm. Mr. Rosenberg is a 20+ year computer industry veteran with technical expertise spanning Local and Wide Area Networking and Internetworking, Host systems integration, software development and integration, Disaster Recovery and Information Security, as well as extensive consulting and project management experience.

Mr. Rosenberg holds a wide range of computer industry certifications, including the following:

- Certified Information Systems Security Professional ("CISSP")
- Symantec Certified Security Practitioner ("SCSP") and Technology Architect ("SCTA") for Vulnerability Management and Firewalls
- Cisco Certified Design Associate ("CCDA")
- Microsoft Certified Professional ("MCP")
- Novell Certified NetWare Engineer ("CNE")

Mr. Rosenberg is a frequent speaker at information security and networking events, and has presented on behalf of Microsoft, LegalTech, the NY State Society of CPAs and other organizations. Mr. Rosenberg has also written articles for the NJ and NY State Society of CPA newsletters and magazines, and authors QTS' QuikNews monthly email newsletter. He has also provided feedback at an executive and a product team level to Microsoft, Novell and other partners.

Mr. Rosenberg has spent the last 14+ years building QTS into one of the industry's leading regional integrators. He has developed numerous product offerings, including the launch of QTS' QuikDesign, QuikSecure and QuikRecover line of offerings, and QTS' service delivery methodologies.

Prior to QTS, Mr. Rosenberg served for seven years at Blue Cross Blue Shield of New Jersey, as an analyst and manager in the Contract Development/Legal Department. Mr. Rosenberg helped modernize and automate the organization's paper-bound workflows, installing the company's first Novell LAN and leading development of its mainframe-based benefits administration system along with numerous other automation initiatives.

Mr. Rosenberg earned a B.A. with Honors in History and English from Rutgers College, and a J.D. degree from Rutgers School of Law - Newark. He passed the New Jersey Bar and served as a member of the NJ State Bar Association's Law Office Management Committee for several years.

Microsoft GOLD CERTIFIED Partner

Citrix access PARTNER GOLD Solution Advisor

symantec. Symantec Gold Partner

Cisco Systems PARTNER PREMIER CERTIFIED

Double-Take Software

SurfControl

hp invent

QTS also offers:

- **QuikSecure** Security Lifecycle services;
- **QuikAssist** support plans;
- **QuikAlert** remote server monitoring;
- **QuikStart** knowledge transfer engagements
- **QuikDesign** planning & architecture services
- **QuikDeploy** Windows desktop deployment services

About Quality Technology Solutions

Quality Technology Solutions is one of the leading network integrators and security service providers in the NJ/NY metro area. QTS has been serving businesses in New Jersey and New York for over 14 years, leveraging a senior talent model to deliver award-winning service and expertise to its customers.

Quality Technology Solutions is:

- Microsoft Gold Certified Partner for Security Solutions, Network Infrastructure, Advanced Infrastructure and Information Worker Solutions;
- Winner of Microsoft Worldwide Partner Awards in 2004 & 2005;
- Microsoft's NJ Medium Enterprise Partner of the Year for 2001-2002, and NYNJ Partner Excellence Award Winner in 2005 and 2006;
- Winner of Novell's Service Excellence Award in 2000 and 2001, one of 15 companies in all of North America to win the award;
- Citrix Gold Partner, Symantec Gold Partner and Cisco Premier Partner;
- Partner with HP, Double-Take Software, Captaris, SurfControl, RSA Security, Good Technology and others to provide a comprehensive range of network solutions.

QTS' goal is to provide its customers with Worry-Free Network Solutions – well-planned and well-implemented projects to ensure networks are reliable, secure and highly available. An important element of this is QTS' QuikSecure methodology, which it utilizes to maximize the Confidentiality, Integrity and Availability of its customers' networks and data resources.

QTS' ***QuikRecover***™ Services presents a complete approach to the lifecycle of implementing a Disaster Recovery Plan for mid-size businesses. Each stage can be a distinct engagement, or the components can be done together in sequence. QTS' QuikRecover solutions include:

- QuikRecover Disaster Recovery Planning Workshop
- Development of System Recovery Plans
- Build-Out of Hot/Warm/Cold Sites
- Disaster Recovery Plan Testing
- Disaster Recovery Plan Review
- Disaster Recovery Plan Documentation

QTS also offers Network Audits, High Availability Reviews and other related services to help you identify and eliminate network weaknesses and reduce the exposure of your business to network outages.

Providing Worry-Free Network Solutions!

Quality Technology Solutions, Inc.
1639 Route 10, Suite 103, Parsippany, NJ 07054
(973)984-7600 www.QTSnet.com