technical® SUPPORT

Supporting Enterprise Networks and Operating Environments

SEPTEMBER 2009 VOLUME 1, NUMBER 2

September is National Preparedness Month.

How Prepared Are YOU?

Part 1 of 2 —What Defines a Next-Generation and Virtual Data Center?

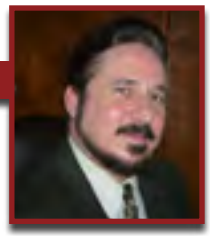

Message from the President

This issue of Technical Support is a little different from those we have published in the past, and in keeping with the notion that as the needs of NaSPA's members change, our organization must also change.

As always, there is a segment of our membership that needs help becoming employed, staying employed or becoming re-employed. In that regard we continue our tips from recruiting expert Jeff Snyder on honing your employment skills and creating your own personal brand.

For those members who really need help, find Part I of a series on considerations before filing bankruptcy. Believe me, we discussed this article for a while as editors. We didn't want to present a "dark" or pessimistic subject. In the end, however, we decided to run the article because financial hardship is a reality for many members, especially with the present economy. We hope this insightful series helps our more needy members understand their options, even in the worst case. In the meantime, have you looked at NaSPA's employment site lately or uploaded your resume?

On the brighter side, we expect to have a NaSPA sponsored link on our web site very soon for members to find health insurance. Watch for it! And speaking of benefits, have you looked at the Liberty Mutual link for your homeowners and car insurance? I have, and the savings are real!

The overall theme of this issue is September 2009, National Preparedness Week. This month is as good time as any to revisit your organization's (and your personal) preparedness for when nature or happenstance throws a curve. There are some great articles inside this edition on being prepared and about Situational Awareness when disaster strikes.

Finally, lest we forget, we have some very cool articles on technology! Ever have to move a data center? Find some great tips here from an expert who has lived through many and can still tell about it. Next, check out Part II of last issue's Performance Optimization in the Virtual Infrastructure Environment series. And to round things out, there is a great piece on electronic reading devices for those of you who, like me, can't resist new gadgets.

I would like to offer my personal thanks to all of the contributors to Technical Support this time around. Because of vacation schedules, summer is a time when it is difficult to nail people down to contribute meaningful articles and content. Not this time! We are delighted as always with the quality of what our authors have delivered. Please let us know how we are doing, and what other articles of interest we can publish, or you can contribute. Email your thoughts to president@naspa.com. Have a safe and happy summer.

Leo A. Wrobel
Editor in Chief Technical Support Magazine
President, NaSPA
President@Naspa.com

technical SUPPORT
Supporting Enterprise Networks and Operating Environments
NaSPA

September 2009
Volume 1, Number 2

Published exclusively for NaSPA, Inc.
Network and Systems Professionals Association

www.NaSPA.com

7044 South 13th Street, Oak Creek, WI 53154
(414) 908-4945, (414) 768-8001 Fax

September 2009 Volume 1, Number 2

FEATURES

ARTICLES

DEPARTMENTS

From the Founder

Waaay back in '68 I learned Fortran IV on an IBM 1401. As I recall I got a D in the course. I couldn't quite understand anything about computing back then. I remember asking my high school Fortran teacher which device was the compiler. You laugh, I know. It's the truth though. Now it's 41 years later and I still write code though now it's SQL database code on a small-company back office system. I've written Fortran, Basic, Autocoder, Fargo, Assembler, Cobol and Rbase SQL. I programmed IBM 1401's, 1620's, 1130's, 360's, 370's, 3033's and that all took place before 1986. Since then I've been totally pc/lan/internet based. I'm still writing code a few hours each week and I still enjoy it very much, though I have many other responsibilities that I enjoy also.

My youngest son Matt graduated in June from high school. He has always enjoyed computers just as much as anyone else at the tender age of 18. Games, pc's, Xbox, Wii and others I can't remember any more. But he now works as a network analyst and is enrolled in college with a major of network administration and he starts college next week.

Perhaps he picked up some of his fondness of computing from his dad, I don't know. I do know two things about this that I'm willing to share. One, my IT career has served me well and lasted a lifetime with challenge, reward, opportunity and fun. I hope that my son gets the same out of it that I do! And two, I'm incredibly proud of him. Imagine, another generation with IT skills. Perhaps the acorn really doesn't fall far from the tree after all.

Best regards,
Scott Sherer
Founder and Past President

NaSPA Mission Statement:

The mission of NaSPA, Inc., a not-for-profit organization, shall be to serve as the means to enhance the status and promote the advancement of all network and systems professionals; nurture member's technical and managerial knowledge and skills; improve member's professional careers through the sharing and dispersing of technical information; promote the profession as a whole; further the understanding of the profession and foster understanding and respect for individuals within it; develop and improve educational standards; and assist in the continuing development of ethical standards for practitioners in the industry.

NaSPA serves Information Systems technical professionals working with z/OS, OS/390, MVS, VM, VSE, Windows Operating Systems, Unix, NetWare and Linux.

September is National Preparedness Month. How Prepared Are YOU?

By Sharon M. Wrobel

NaSPA Board Member Raymond V. Hall, Training and Consultant Services, dropped us an e-mail last month reminding us that September is National Preparedness Month (NPM). This was a very timely reminder that struck a chord with me, NaSPA member Sharon Wrobel. You see I am starting to get active in NPM through our local Homeland Security/ State Civil Defense organizations as well as Community Emergency Response Teams or CERT. It's partly because I work in disaster recovery and partly because I have a desire to give back to the community.

While preparedness for disasters is an all-year long endeavor, NPM affords us all the opportunity to ask ourselves how well WE are prepared, not only as organizations but as individuals.

For the sixth year in a row, NPM will be held in September. It is designed to encourage Americans to take the simple steps to prepare for emergencies in their homes, businesses, and communities.

According to NPM, National Preparedness Month Coalition membership is open to all public and private sector organizations. Groups can register to become a NPM Coalition Member by visiting www.ready.gov and clicking on the NPM banner.

During NPM, Coalition Members share preparedness information with their members, customers, employees and communities. In fact, last year, the Ready Campaign partnered with more than 3,200 organizations to highlight the importance of public emergency preparedness throughout September. This year the focus will be on changing the way one thinks about emergency preparedness and helping others understand what is meant to be READY.

Throughout the year, the *Ready* Campaign promotes individual emergency preparedness. A national advertising campaign, produced in partnership with The Advertising Council, is designed to educate and empower Americans to prepare and respond to emergencies, including natural disasters, terrorist attacks and other threats.

The Campaign lists its web sites as www.ready.gov and www.listo.gov; with two toll free numbers, 1-800-BE-READY and 1-888-SE-LISTO that are also available to provide Americans with free emergency preparedness information.

So what are some of the things that you can do as an individual or business to promote National Preparedness Month?

1. You can do as I just did and write an article that will reach thousands of residents, employees, or members of an organization such as NaSPA, the Chamber of Commerce or even in your corporation's newsletter.
2. Enlist your local fire or police department to give instruction in CPR, Emergency training or First Aid or hold a blood drive.
3. Join one of the organizations below where you will receive instruction on how to help yourself and others and where you will be able to participate in mock disaster exercises.
4. Utilize your talents and education by assisting others if you are in one of the service oriented professions such as doctor, nurse, or veterinarian.

If you are so inclined to join, the following contains a few of the names of organizations dedicated to promoting public awareness

Citizens Corps

The mission of Citizen Corps is to harness the power of every individual through education, training, and volunteer service. They pledge to make communities safer, stronger, and better prepared to respond to the threats of terrorism, crime, public health issues, and disasters of all kinds. Citizen Corps promotes a range of measures for you to make your family, your home, and your community safer from the threats of crime, terrorism, and disasters of all kinds.

through education, training and volunteer service. What you will find is that each and every one of these fine organizations has similar traits and goals.

Local Fire and Police Departments offer classes to citizens interested in learning more about personnel and policies and taking an active role in helping to protect their communities. Class topics may include:

1. Emergency medical training,
2. Communications,
3. Codes and
4. Hands-on training.

As stated in its web site, Citizen Corps brings together community and government leaders to tackle issues at the local level including emergency preparedness, planning, mitigation, response and recovery. These Councils enable collaborative planning between government and civic leaders and provide localized support for outreach and educational efforts directed to the public. Training and exercises effectively integrate emergency responders, volunteers and the general public that encompass the full range of emergency response services.

To find an academy near you, contact your local police or fire department. If you'd like to take an active role in helping your community prepare for emergencies, join Citizen Corps. Get information at www.citizencorps.gov

Community Emergency Response Teams (CERT)

The Community Emergency Response Team (CERT) Program can be found at www.citizencorps.gov/cert/. It provides education and prepares citizens on how to respond to potential disaster hazards that may have an effect on their area by training them in basic disaster response skills. These might include such skills as fire safety, light search and rescue, team organization, and disaster medical operations. Following an event, CERT members apply what they have learned in the classroom and training exercises by aiding residents and employees sometimes even before emergency responders have had a chance to arrive. By taking a more active role in emergency preparedness projects in their community, CERT members become a support system alongside emergency response agencies.

Cities Readiness Initiative (CRI)

Federally funded, the Cities Readiness Initiative (CRI) prepares major US cities and metropolitan areas to effectively respond to a large scale bio-terrorist event by dispensing antibiotics to their entire identified population within 48 hours of the decision to do so.

I recently had the opportunity to participate in one such exercise. CRI along with the Texas State Civil Defense staged an anthrax event in a nearby city. I, being a "client" who wasn't sick yet, had to wait in line with other clients for my turn, be assessed as to how sick I was then either brought to the next station or to a holding room. Security then checked my identity and moved me along to the processing station where I filled out paperwork. I was then moved to the dispensing

station and given my anthrax medication (tic tacs in this case) which ended my participation. Everything had to move smoothly and the people assisting had to know what they were doing at their particular station.

CRI's success comes from enhanced communication and collaboration across state and local boundaries, resulting in optimal use of shared resources. Through exercises, assessments and reviews, local and state planners have identified capability, strengths, and shortcomings. Success is also hinged upon the availability of federal resources to local areas.

For more information consult the CRI web site www.bt.cdc.gov/cri/

"The President proposes to create a new Department of Homeland Security, the most significant transformation of the U.S. government in over half-century by largely transforming and realigning the current confusing patchwork of government activities into a single department whose primary mission is to protect our homeland. The creation of a Department of Homeland Security is one more key step in the President's national strategy for homeland security."

-From the Department of Homeland Security

June 2002 - George W. Bush

Department of Homeland Security

It's not surprising to anyone that after the bombing of the World Trade Center, September 11, 2001, our country was forced into National Preparedness. No longer were we to be complacent. No longer did we, in the United States of America, think we would be spared from evil because after all, we were the good guys; protector of other nations. And so, the Department of Homeland Security was born to provide a terrified nation a security blanket that would not tolerate another terrorist attack but would take on the primary responsibility of educating emergency responders so that they would be able to take control, respond and assist with recovery www.dhs.gov/index.shtm

American Red Cross

One can not talk about preparedness without giving due credit to the first organization that offered emergency response to our nation. Founded in 1881 by visionary leader Clara Barton, the American Red Cross differentiates itself by coming to the aid of victims of natural or manmade disasters. They are the first ones to show up, carrying with them all the comforts and compassion that a family numb with disbelief needs right at that moment. Although American is in their name, they venture worldwide providing unbiased care to those in need. Their main goal is to provide relief and prevent suffering.www.redcross.org/

If you are still asking yourself **ARE YOU READY?** take a survey from the Red Cross at www2.redcross.org/preparedness/npm/ entitled, *What's Your Readiness Quotient?*

September is National Preparedness month. I challenge all readers and members to make preparedness an action item. How will YOU prepare?

Sharon M. (Ford) Wrobel is managing editor of Technical Support for NaSPA. Sharon conducted extensive publishing and regulatory research for her former employer, (a 50 state telephone company), a function she continues today as Vice President of Business Development for b4Ci Inc. Sharon was a major content contributor to NaSPA's president Leo A. Wrobel's 2009 book "Business Resumption Planning Second Edition" © Taylor Publishing Inc, and was co-author of Leo's latest book "Disaster Recovery for Communications and Critical Infrastructure" © 2009 Artech House Books Inc. She can be reached at sharon@b4ci.com.

Part 1 of 2 —What Defines a Next-Generation and Virtual Data Center?

By Greg Schultz

Virtual data centers are increasingly well known to enable data mobility, resiliency, and improved IT efficiency. In this helpful series, based on the book by Greg Schultz *"The Green and Virtual Data Center"* you will learn:

- What defines a green and virtual data center
- How virtualization can be applied to servers, storage, and networks
- The many faces of server, storage, and networking virtualization
- How to leverage virtualization beyond server, storage, and network consolidation
- The various components and capabilities that comprise a virtual data center
- How existing data centers can transform into next-generation virtual data centers

Special thanks to NaSPA supporter Auerbach Publishing, as well as to Greg Schultz for their permission to use this timely and useful material.

Considering Going "Green?"

Many approaches and technologies, addressing different issues and requirements, can be used to enable a green and virtual data center. Virtualization is a popular approach to consolidating underutilized IT resources, including servers, storage, and input/output (I/O) networks to free up floor-space, lower energy consumption, and reduce cooling demand, all of which can result in cost savings. However, virtualization, and particularly consolidation-applies to only a small percentage of all IT resources. The genesis of this article series is that there are many facets of virtualization that can be used to enable IT infrastructure resource management to improve service delivery in a more cost-effective and environmentally friendly manner.

In my book, "Virtual Data Center" we discuss various green and environmental issues. In this article, based on that book, we turn to what can be done and how to leverage different technologies to enable a green virtual data center. This series also looks at how virtualization can be applied to servers, storage, and networks. Virtualization can be applied to consolidate underutilized IT resources. It can also be used to support transparent management, enabling scaling for high growth and performance applications in conjunction with clustering and other technologies.

Why Virtualize a Data Center?

A virtual data center can, and should, be thought of as an information factory that needs to run 24-7, 365 days a year, to deliver a sustained stream of useful information. For the information factory to operate efficiently, it needs to be taken care of and seen as a key corporate asset. Seen as an asset, the IT factory can be invested in to maintain and enhance productivity and efficiency, rather than being considered a cost center or liability.

The primary focus of enabling virtualization technologies across different IT resources is to boost overall effectiveness while improving application service delivery (performance, availability, responsiveness, security) to sustain business growth in an economic and environmentally friendly manner. That is, most organizations do not have the luxury, time, or budget to deploy virtualization or other green-related technologies and techniques simply for environmental reasons-there has to be a business case or justification.

Virtual data centers (See Figure 1), regardless of whether new or existing; require physical resources including servers, storage, and networking, and facilities to support a diverse and growing set of application capabilities while sustaining business growth. In addition to sustaining business growth, applications need to be continually enhanced to accommodate changing business rules and enhance service delivery. Application enhancements include ease of use, user interfaces, rich media (graphics and video, audio and intuitive help), along with capturing, storing, and processing more data.

There are a growing number of business cases and justifications for adopting green technologies that reduce costs or maximize use of exist-

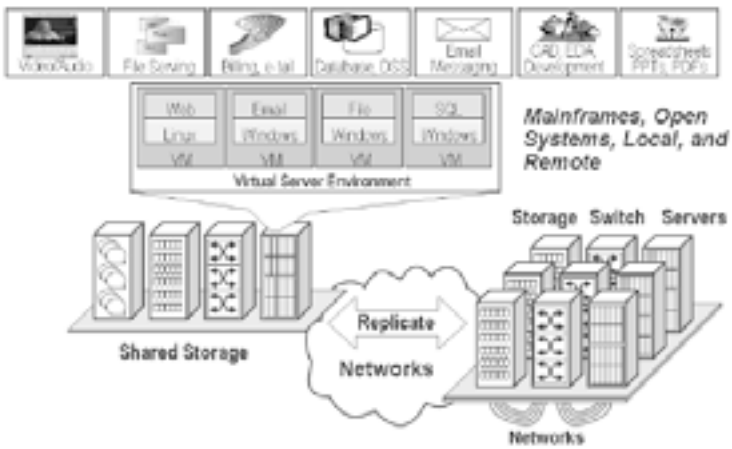

Figure 1. Virtual Data Center with Physical Resources

ing resources while also benefiting the environment. However, it is rare for a business to have surplus budget dollars, personnel, facilities, and management support to deploy new technologies simply for the sake of deploying them. Challenges facing IT data centers and businesses of all sizes include:

- Ensuring reliability, availability, serviceability, and management
- Dealing with performance bottlenecks and changing workloads
- Securing applications and data from diverse internal and external threat risks
- Enabling new application features and extensibility while supporting business growth
- Managing more data for longer periods of time
- Supporting compliance or self-governance while securing intellectual property data
- Maximizing **p**ower, **c**ooling, **f**loor space, and **e**nvironmental (PCFE) issues
- Ensuring data and transaction integrity for business continuance and disaster recovery
- Consolidating IT resources to reduce complexity and contain management costs
- Supporting interoperability with existing applications, servers, storage, and networks
- Reducing or abstracting complexities associated with applications interdependencies
- Leveraging existing IT personal skill sets and experience
- Scaling existing and new applications with stability and in cost effectively
- Moving data transparently from old to new storage during technology replacement
- Enhancing IT service delivery in a cost-effective and environmentally friendly manner

IT data centers are increasingly looking to virtualization technology and techniques to address all these issues and activities. Virtualization can be used to consolidate servers and storage resources in many environments to boost resource utilization and contain costs-for example, using a software based virtual machine (VM) to combine applications and operating systems images from underutilized physical servers to virtual machines on a single server. With a growing focus on PCFE and the greening of IT in general, consolidating servers, as well as storage, networking, and facilities, is a popular and easy-to-understand use of virtualization.

Energy efficiency today can sometimes mean simply energy avoidance. In the future, however, emphasis will shift to doing and enabling more work and storing increasing amounts of information for longer periods of time. This will have to be accomplished while consuming less energy and using less floor space. Doing more work in a more productive manner with less energy will result in efficiencies from improved technologies, techniques and best practices.

Consequently, a green and virtual data center should be much more than just an environment that leverages some virtualization for consolidation purposes. A green and virtual data center should enable transparent management of different physical resources to support flexible IT service delivery in an environmentally and economically friendly manner on both local and remote bases.

Scaling with stability means that as performance is increased, application availability or capacity is neither decreased nor is additional management complexity or cost introduced. Similarly, scaling with stability means that as capacity is increased, neither performance nor availability suffer, nor is performance negatively affected by growth or increased workload or application functionality. This also includes eliminating single points of failure and supporting fault isolation and containment, self-healing, supporting mixed performance of small and large I/O operations, and additional functionality or intelligence in technology solutions without adding cost or complexity. In addition, scaling with stability means not introducing downtime or loss of availability or negative performance as a result of growth. High resilience or self-healing with fault isolation and containment will prevent single points of failure from cascading into rolling disasters.

Flexibility and agility will enable virtual data centers to meet changing business and application requirements while quickly and transparently adopting new technology enhancements without disruption.

Tenets of a green and virtualized data center include:

- Flexible, scalable, stable, agile, and highly resilient or self-healing systems
- Quick adaptation and leverage of technology improvements
- Transparency of applications and data from physical resources
- Efficient operation without loss of performance or increased cost complexity
- Environmentally friendly and energy efficient yet economical to maintain
- Highly automated and seen as information factories as opposed to cost centers
- Measureable with metrics and reporting to gauge relative effectiveness
- Secure from various threat risks without impeding productivity

Virtualization Beyond Consolidation-Enabling

Transparency

There are many facets of virtualization. Aggregation is a popular approach to consolidate underutilized IT resources including servers, storage, and networks. The benefits of consolidation include improved efficiency by eliminating underutilized servers or storage to reduce electrical power, cooling, and floor space requirements as well as management time, or to reuse and repurpose servers that have become surplus to enable growth or support new application capabilities.

Another form of virtualization is emulation or transparency providing abstraction to support integration and interoperability with new technologies while preserving existing technology investments and not disrupting software procedures and policies. Virtual tape libraries are a commonly deployed example of storage technology that combines emulation of existing tape drives and tape libraries with disk-based technologies. The value proposition of virtual tape and disk libraries is to coexist with existing backup software and procedures while enabling new technology to be introduced.

Figure 2 shows two examples of virtualization being used, with consolidation on the left side and transparency for emulation and abstraction to support scaling on the right. On the consolidation side, the operating systems and applications of multiple underutilized physical servers are consolidated onto a single or, for redundancy, multiple serv-

ers in a virtual environment with a separate virtual machine emulating a physical machine.

In this example, each of the operating systems and applications that were previously running on their own dedicated servers now run on a virtual server. Consolidation enables multiple underutilized servers to be combined yet let each system think and operate as though it still had its own server.

For a variety of reasons, not all servers or other IT resources lend themselves to consolidation. These reasons may include performance, politics, finances, service-level, or security issues. For example, an application may need to run on a server with low CPU utilization to meet performance and response-time objectives or to support seasonal workload adjustments. Also, certain applications, data, or even users of servers may need to be isolated from each other for security and privacy reasons.

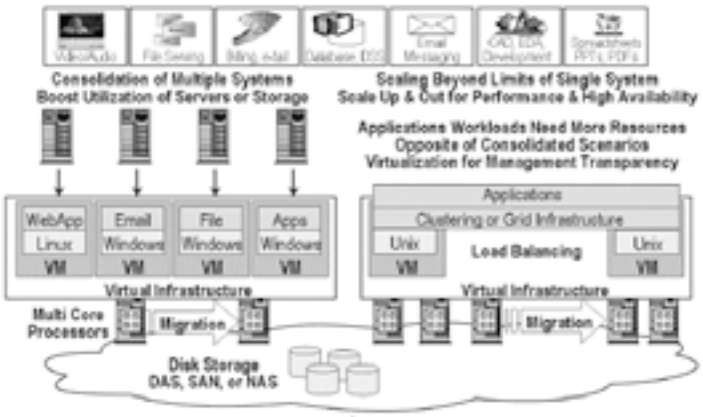

Figure 2. Consolidation vs. transparent management with virtualization

Politics and financial, legal, or regulatory requirements also need to be considered. For example, a server and application may be "owned" by different departments or groups and thus managed and maintained separately. Regulatory or legal requirements may dictate that certain systems be kept separate from other general-purpose or mainstream applications, servers, and storage. Separation of applications may also be necessary to isolate development, test, quality assurance, back-office, or other functions from production or online applications and systems, as well as to support business continuance, disaster recovery, and security.

For applications and data that do not lend themselves to consolidation, a different way to use virtualization is to enable transparency of physical resources to support interoperability and coexistence between new and existing software tools, servers, storage, and networking technologies, such as enabling new, more energy-efficient servers or storage with improved performance to coexist with existing resources and applications.

On the right side in Figure 2 are examples of applications that don't lend themselves to consolidation because they need to be isolated from other applications or clients for performance or other reasons. However, these applications can still benefit from transparency and abstraction when combined with clustering technology to enable scaling beyond the limits of a single server, storage, or networking device. Included on the right side of Figure 2 are applications that need more resources for performance or availability than are available from a single large server and hence need to scale horizontally (also known as scale up and out).

Another use of virtualization transparency is to enable new technologies to be moved into and out of running or active production environments to facilitate technology upgrades and replacements. Still another use is to adjust physical resources to changing application demands, such as seasonal planned or unplanned workload increases. Transparency via virtualization also enables routine planned and unplanned maintenance functions to be performed on IT resources without disrupting applications and users of IT services.

Virtualization in the form of transparency or abstraction of physical resources to applications can also be used to help achieve energy savings and address other green issues by enabling newer, more efficient technologies to be adopted more quickly. Transparency can also be used to implement tiered servers and storage to leverage the right technology and resource for the task at hand as of a particular point in time.

Business continuity (BC) and disaster recovery (DR) are other areas in which transparency via virtualization can be applied to in a timely and cost-efficient manner in-house, via a managed service provider, or using some combination. For example, traditionally, a BC or DR plan requires the availability of similar server hardware at a secondary site. Some challenges with this kind of redundancy are that the service and servers must be available to those who need them when they need them. For planned testing, this may not be a problem; however, in the event of a real disaster, a first-come, first-served situation could arise, with too many subscribers to the same finite set of physical servers, storage, and networking facilities.

If dedicated and guaranteed servers and storage resources are available for BC and DR, competition for resources is eliminated. This means, however, that additional servers and storage need to be powered, cooled, and given floor space and management. In addition, these operating systems and applications may require identical or very similar hardware and configurations.

Part 1 Summary

For many organizations, the opportunity to start from scratch with a nice, new green data center may exist. However, for most of us, enabling a virtualized data center relies heavily on transitioning existing facilities, servers, storage, networking, and software tools, processes and procedures to a virtualized environment. (See our "Smooth Moves" article in this month's *Technical Support Magazine* by Phil Rothstein for other useful tips on data center transitions!) The benefits are many and include not only a "green" mentality toward your data center but also support, growth and enhancement of service delivery in a cost-effective and timely manner. It's worth the trouble. More tips next time in Part 2 of this series, or if you just can't wait, pick up my book from Auerbach and learn more! To order click HERE.

Part 1 of 2—What Technologists Should Know Before Considering Bankruptcy

By Billy D. Price, Attorney

The editorial staff here at *Technical Support* will admit that we found this an unsettling topic when it was first proposed. We decided however that the topic is not only timely, it's an 800 pound gorilla for many technologists during these trying economic times, and one that bears some honest discussion.

If you are out of work, or have determined that you cannot reorganize your finances, you may be considering filing for bankruptcy.

It's no secret these days that many professionals, NaSPA members included, are facing financial difficulty and an uncertain financial future. While all of us experience financial pressure from time to time, these are indeed tough times for many. The financial crisis, the credit crisis, mortgage crisis, 401K's and stocks going up in smoke. Downsizing, rightsizing and capsizing. Escalating health care costs. Outsourcing and exporting jobs. For many it's like the old Chinese saying, "death by 1000 cuts."

If you are out of work, or have determined that you cannot reorganize your finances, you may be considering filing for bankruptcy. If so, this is likely to be an emotional time for you. That's why we decided to run this two part series. In addition to the other initiatives we are taking to help NaSPA members during these trying economic times, (like insurance programs and job placement services) we also want NaSPA members in financial trouble to be aware of all the facts before making permanent decisions. That's why we approached an attorney to write this series and provide a confidential resource to NaSPA members in financial trouble. If you or someone you know is considering bankruptcy, you may be interested in some of the tips in this article. If you are fortunate enough to be weathering the present recession well, perhaps you can pass this on to someone less fortunate. If you would like some quality time with Billy Price, the author of this series, to discuss your options, contact him with your NaSPA membership number. There will be no charge for the initial consultation and we have arranged for a NaSPA discount ($250.00) if you find you must file. His number is at the end of this article.

What is Bankruptcy?

It's easy to trivialize bankruptcy and attribute it to overspending. But more often than not, most everyone has probably lived through an event that has temporarily affected their income or permanently altered their standard of living. One may have lost a job, divorced a spouse, lost a loved one, lost a business, or incurred substantial medical expenses. In this article we will attempt to answer common bankruptcy questions.

What happens when someone files bankruptcy?

When someone files bankruptcy, his or her case is assigned to a Bankruptcy Trustee. The Bankruptcy Trustee is a neutral party who determines how to divide the petitioner's assets in order to pay creditors. The trustee decides which debts will be paid completely, partially, or not at all. Because of the nature of bankruptcy, bankruptcy attorneys can predict which debts must be paid, and which debts may be discharged or written off. The type of bankruptcy filed is determined by many of these factors. The different types of bankruptcies are described below.

What is the court's role in a bankruptcy proceeding?

The court is the ultimate decision maker in deciding whether to approve or reject the terms of the bankruptcy case. Except in rare cases, the Bankruptcy Court approves most bankruptcy cases.

Does a Petitioner have to go to court?

In most bankruptcy cases, a petitioner for bankruptcy will only attend a "meeting of creditors" in which they and their attorney meet with the Bankruptcy Trustee and any creditor who attends. Usually, creditors choose not to attend. Typically, this is a short meeting and the petitioner is asked some simple questions concerning their bankruptcy forms and financial situation. Occasionally, if complications arise or a petitioner chooses to dispute a debt, the petitioner may be required to appear before a judge at a hearing. If they need to go to court, they will receive notice of the court date and time from the court and their lawyer.

What are the ultimate consequences of filing bankruptcy?

After someone files bankruptcy, their creditors can no longer attempt

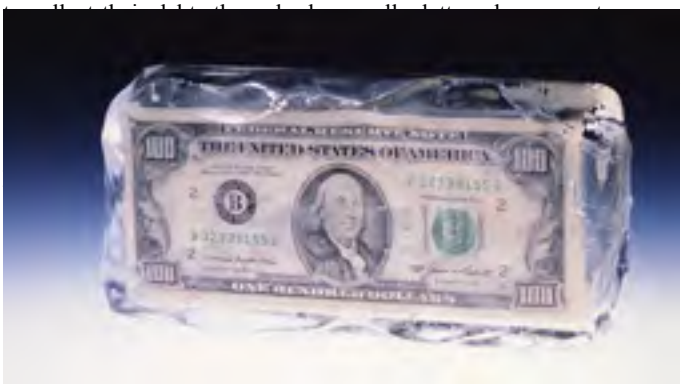

is not secured by collateral. For example, if a petitioner owes car payments, the car is the collateral for the loan on the car and therefore, a secured debt. In contrast, there is no collateral in an unsecured debt. Unsecured debt includes, but is not limited to 1) credit card debt; 2) medical bills; 3) old accounts; 4) old bills; 5) charge-offs; and 6) repossessions and foreclosures. A lawyer can help determine which debts are secured and which are unsecured.

Are there different forms of bankruptcy?

Yes. The Bankruptcy Code provides bankruptcy solutions for individuals, family farmers and fishermen, businesses, and cities. The most important bankruptcy solutions for individual debtors are a Chapter 7 Bankruptcy and a Chapter 13 Bankruptcy, both of which are described in more detail below.

What is the definition of an individual under the Bankruptcy Code?

Under the Bankruptcy Code, an individual includes single people and married couples. Married couples are considered "individuals" because married couples often share debts. Typically, married couples file bankruptcy together. If debt is shared with someone other than a spouse, a petitioner may not file a joint bankruptcy with that person. For example, imagine that a petitioner's mother is the cosigner on the loan for that person's house. Even though the petitioners and their mother share the debt, the petitioner's mother cannot file a joint bankruptcy with them. To obtain debt relief, the mother must file bankruptcy herself.

Can Property Be Lost?

Determining whether a petitioner can retain property depends on the type of bankruptcy filed (which is described in more detail below) and whether money is owed on the property. If a petitioner owns property outright, *exemptions* may enable them to keep the property.

What is an exemption?

Exempt property is property which is owned outright and which may be kept despite filing bankruptcy. The best way to explain an exemption is to describe what would happen if the law did not allow for exemptions. For instance, pretend that a friend of yours owns all of the furniture in her home outright (this means that her furniture is not collateral for any debt. In other words, her furniture is completely "paid off") and that she also owes unrelated credit card debts. If she were not allowed to claim exemptions, her unsecured creditors could force her to sell property that she has already paid off in order to pay her other current debts. To elaborate, continue with our previous hypothetical in which your friend owns her furniture outright, but also owes credit card debt. Without exemptions, her credit card company could force her to sell her furniture to pay her credit card debt. That does not typically happen because of exemptions.

Exemptions allow the petitioner to keep a certain amount of property which is owned outright without having to worry that other creditors can claim an interest in that property. Most petitioners are able to exempt all the property they own outright.

When should someone consider bankruptcy?

Filing for Bankruptcy is a very important decision and should be carefully weighed in terms of the pros and cons. A bankruptcy can be advantageous in the following circumstances, if a person:

1. has foreclosures or repossessions on their credit;
2. cannot pay the minimum on credit card bills each month;
3. finds that minimum monthly credit card payment does not reduce the balance owed;
4. has more medical bills than they can afford;
5. is receiving harassing phone calls from creditors and collection agencies;
6. has a low credit score and is unable to purchase items on credit or open a new line of credit; or
7. has been sued by one of their creditors

Sometimes, a potential client comes to me and after discussion we decide that a bankruptcy is not the best choice for him or her. However, I am nearly always able to assist the people who request my help. Another good resource is the Consumer Credit Counseling Service (CCCS). CCCS is not affiliated with any legal service or attorney. CCCS is a U.S. Trustee approved provider of pre-filing counseling and pre-discharge education. This is important for two reasons. First, these folks can help you go in with both eyes open and fully aware of your options through unbiased, objective information, education, and counseling. For example, the new Bankruptcy laws enacted in 2005 encourage a petitioner to file a Chapter 13 (which requires a partial payback of debt) rather than a Chapter 7 (which writes ALL of it off) if their income is above the median income for the area in which they live.

An organization like CCCS can "pre qualify" a potential petitioner and provide them with what their options are. A firm of this type might help a petitioner avoid bankruptcy altogether, by acting as an intermediary with creditors and negotiating a *debt settlement* with creditors instead of a bankruptcy.

The second reason CCCS is a good idea is that the new Bankruptcy laws *require* this kind of counseling as a precondition of a bankruptcy being approved. A petitioner must complete two separate sessions of financial counseling of this type before their filing can be ultimately approved.

Will bankruptcy affect a Petitioner's credit?

There is no clear answer to this question. Unfortunately, if a person is behind on their bills then their credit may already reflect this. If someone files bankruptcy, it is unlikely that his or her credit will worsen. The bankruptcy will, however, appear on a person's credit report for 10 years if they file a Chapter 7 Bankruptcy and 7 years if they file a Chapter 13 Bankruptcy. Bankruptcy wipes away old debts however. It is likely therefore that the petitioner will be better equipped to pay their current debts, and in fact, may be able to obtain a new line of credit because of that fact.

Should a Petitioner hire an attorney to represent them in a bankruptcy proceeding?

Although it may be possible to file bankruptcy without an attorney's assistance, this may be an unwise decision. The bankruptcy process includes many legal technicalities and the person considering it must remember that their property is at stake. If they do not fully understand the law and their legal rights, it is more likely that they will lose property, or all of their debts will not be discharged. Filing a bankruptcy requires patience and careful preparation. Sometimes, people are able to file a Chapter 7 Liquidation Bankruptcy without an attorney, but a

Although it may be possible to file bankruptcy without an attorney's assistance, this may be an unwise decision.

Chapter 13 Reorganization Bankruptcy is more difficult. Few people have been able to complete a Chapter 13 Bankruptcy on their own.

Article Summary

We hope the first part of this series has answered some Bankruptcy Questions. We will conclude the series in the next series in greater detail concerning the law and a petitioner's options. Remember, however, the best article is not a substitute for a meeting with an attorney who can discuss your finances and options in confidence. If you would like to set up a visit or a conference call, please call (214) 696-9601. Consultation is free, and a discount is available with your NaSPA member number if filing is appropriate. If you have additional questions, you may send you e-mail questions to the author in complete confidentiality at billy@billyprice.net.

The Electronic Reading Device Line Up

By Bill Elder

COOL-ER™ ebook device

In the last issue of *Technical Support*, we introduced you to the Kindle electronic reading device and provided an overview. The world of electronic reading devices has gotten to be very competitive. It reminds me of the early years of the Internet Boom. Back then, search engines were relatively new. No one really knew who was going to dominate the search market. The competition is still going on; especially with the development of specialized search engines for just about any subject under the sun.

There seems to be a similar event going on with the electronic publishing industry. In addition to Amazon, many other companies have joined the race to produce the best electronic reading device. Only time will tell who will win this competition for dominance.

As we cover the theme of National Preparedness Week, it's an appropriate time to explore the role of portable devices in dealing with the unexpected. Cell phones have definitely played a part in responding to disasters and saving lives. There are cases where businesses would not have survived without being able to access their key business data through laptops. Electronic reading devices are probably going to play a part in helping to deal with disasters as well

In this issue, we will provide a broad overview of some the main competitors in this industry. We also have an interview with Phil Wood of Interead, a company that produces the COOL-ER™ ebook device. Phil will also provide us with some interesting insights about the possible impact of the electronic reading device on how we as IT professionals will perform our jobs. He also provides some insights about how organizations are starting to look to electronic reading devices for disaster planning.

Keeping up with the Industry

For those who want immediate updates to this fast changing arena, go to the website www.ebookrumors.com. This site posts constant updates and has links to areas such as affordable ebooks, free ebooks, and ebook software. It is a forum for readers, writers and developers involved with ebooks.

The line up of Electronic Reading Device Companies

Since I have been researching this topic, I have been noticing more articles about these devices lately. The main companies getting coverage include these four outfits:

- Amazon (Amazon Kindle 2 and Amazon Kindle DX) www.amazon.com
- Sony (Sony Reader PRS-700) www.sonystyle.com/Reader
- Plastic Logic (eReader Device) www.plasticlogic.com/
- Interead (COOL-ER™ Device) www.coolreaders.com

Developments at Amazon

Amazon recently decreased the price of its Kindle device to $299. The device can hold over 1,500 books and according to the website, readers have "Over 300,000 books plus U.S. and international newspapers, magazines, and blogs available." Amazon is also coming out with the Kindle DX that will also enable readers to hold up to 3,500 books and be able to read online newspapers and magazines such as the New York Times and the Wall Street Journal. This device goes for $489.

The Magazine, Fast Company, recently had an in-depth interview with Amazon's CEO, Jeff Bezos in the July/August 2009 edition titled "The Evolution of Amazon". Jeff Bezos provides similarities between the launch of the Kindle and the launch of the Apple iPOD device. The publishing industry is certaintly going through an upheaval. There does seem to be some parallel to the music industry when it had to adapt to the Napster phenomenon and then the introduction of the iPod and MP3 players. The book and magazine worlds are going through similar fates.

Developments at Sony

If you do a search on the Sony e-reader on sites like Yahoo and Google, you will find many sites that compare the Sony to the Kindle. Even YouTube.com has video demonstrations of these two devices. One example of the strong competition is a May 2009 article, "Amazon vs. Sony: E-Book Reader Face-Off" found on PCWORLD.COM. The Sony and Kindle devices are close in price but have some key differences. The Sony device is more focused on reading and does not include a keyboard which the Kindle does have as a feature. The Sony

Racing menu on COOL-ER™

device only works with PCs with its software package. For those who consult Consumer Reports.com, there are plenty of reviews; particularly of the Sony and Amazon devices.

The latest with Plastic Logic

Plastic Logic recently made a deal with AT&T to make their reading device available through this wireless network. This deal supports another one Plastic Logic made with the Barnes and Noble bookstore. An interesting feature of being part of the AT&T network is that you can use their reading device internationally. Currently, this is not possible through the Kindle. That may change over time. You can read more about these deals at the Plastic Logic website containing the Forbes.com article titled "New e-book reader to use AT&T Network." (07/22/09)

News from Interead

A relative newcomer to the electronic reading device world is a company called Interead with offices in the United States and Great Britan that has introduced a device called the COOL-ER™. A British website, www.seenit.co.uk/, recently had a post featuring this device. One interesting note about this device is that it can handle many different formats such as PDF, EPUB, FB2, RTF, TXT, HTML, PRC, JPG, BMP, GIF.

We contacted Phil Wood, Interead's marketing director, to tell us more about the COOL-ER™ and his insights about the electronic reading device industry.

1. Technical Support-Being in the Information Technology (IT) profession comes with having to buy many books to stay current. In your opinion, how will the electronic reading device industry impact the way IT professionals stay current in the field?

   Interead-The impact will be huge because of freshness and capacity. The most popular and up-to-date IT books are now available to buy and download via the Internet. For the IT profession this is a key fact as the industry changes so quickly you need to be updated about your profession more frequently. The electronic books that are downloadable come in many formats, such as pdf, prc, txt, epub etc. The most popular up and coming format is the epub (epublishing) format. The key benefit of an electronic book reader is that it enables hundreds of generic books/manuals to be stored in one device, all completely fresh and up-to-date.
2. Technical Support-Many of our readers are involved with disaster recovery planning and responding. Being able to respond to disasters often involves the use of portable devices. Have you seen any use of electronic reading devices in response to disasters? How could IT professionals consider using electronic reading devices to help with disaster planning? For example, could these devices be used as portable electronic corporate repositories?
   Interead-We have had interest from the UK government from specific departments, like the National Health Service and Army, for a portable electronic device that can store huge, read-only documents such as emergency procedures and instruction manuals. The need for a fixed, secure, format is key to this. Specifically the COOL-ER™ ebook device uses Digital Rights Management (DRM) from Adobe Digital Editions software to enable a secure operating system to be in place for such emergency situations.
3. Technical Support-Some of our readers also write books and articles. For those who want to publish for the electronic reading device audience, how could your firm play a part in that effort? Any insights to offer about writing and formatting specifically for electronic reading devices?

   Interead-In Q409 Interead plans to have a coolerwriting.com website that will be designed for budding authors and smaller publishers (who find the large data aggregators and publishers too expensive to work with) to directly upload their work. We also plan to have an open platform that allows COOL-ER™ ebook users to upload and share applets that they have designed/written themselves. We believe we will offer a cost effective and direct route to market for what has been a very price prohibitive industry to date.
4. Technical Support-Do you have any additional insights to offer about the way electronic reading devices could impact/change the way IT professionals will perform their jobs?

   Interead-In early 2010 Interead plans to have a roadmap of ebook readers that have a larger, colour screen, with wireless features to enable magazines and the like to be downloadable on the move. Imagine getting your Computer publication to read the night before you can buy/read it via the newsagent?
5. Technical Support-Some of your other competitors include Amazon, Sony, and Plastic Logic. What special features/services does your firm offer that the others may not be able to offer?

   Interead-We believe in openness, so our ebook website, www.coolerbooks.com offers all ebook readers, (except Kindle due to its unique non sharable IP and which is not available outside of the USA anyway) the choice of downloading their devices ebooks. Interead also gives back to the community. We are donating a percentage of our net profits to Literacy Trusts across both sides of the pond.

Side view of COOL-ER™

Interead's company overview is below;

- Company name: Interead- Website url COOL-ER™ ebook device is; www.coolreaders.com/ Our accompanying ebook website is at; www.coolerbooks.co.uk/ - Model name COOL-ER™
- Brief product description: COOL-ER™ is the ebook reader designed to truly fit the modern lifestyle. It's the smallest,

lightest, fastest and most affordable ebook reader ever. Check it out at www.coolreaders.co.uk . It's like having an entire library in one gorgeous electronic package. COOL-ER™ is also the first e-reader to work with books in any format, from any store, and has a companion ebook store at www.coolreaders.co.uk, where books can be bought, downloaded and—best of all—shared.

COOL-ER™ uses an amazing new technology that gives you a clear "like paper" electronic display. COOL-ER™ displays a crisp, clear image on a large screen, replicating the experience of reading a paperback book. The 6 inch display uses power only for page turns, and a single charge lasts around 8,000 page turns, which is like reading 12 average books without re-charging! And unlike a book, you can change the display text; the days of straining your eyes are now gone forever. You can store hundreds of books on the 1GB internal memory. The COOL-ER™ is Windows PC and Mac compatible and supports Adobe Digital Editions. It's 45% lighter than other ereaders.

- Geographic availability: Worldwide
- Price: 225 Euros for Europe £189 inc. VAT & Delivery; $249 in USA & RoW
- Supported eBook file types: (PDF, EPUB) PDF, EPUB, FB2, RTF, TXT, HTML, PRC, JPG, BMP, GIF.
- Device ship date Now shipping Worldwide

How many parents will be pressured to buy these type of reading devices for their kids along with laptops, cell phones, iPods, and much more?

Closing Thoughts

Another device hits the IT world and the world at large. Now that I am writing about these readers, I will be noticing if more people are using them while I wait for flights at airports. Another big question is, how many parents will be pressured to buy these type of reading devices for their kids along with laptops, cell phones, iPods, and much more?

In the not too distant future, we may see some other interesting changes. New generations of IT workers will have been brought up on these reader devices and possibly bring in new ways of doing business. These new ideas might lead to even more innovation in helping all of us prepare for the unexpected.

In our next issue of *Technical Support*, we will continue to provide industry overviews and interview other companies involved with these dynamic gadgets. Please stay tuned!

NaSPA Takes It on the Road!

NaSPA staffed a booth at the NDR (National Disaster Recovery) show in New Orleans last month. Pictured are Leo and Sharon Wrobel, and other NaSPA friends. NaSPA members earned free admission to NDR by virtue of their NaSPA membership and the booth was a big hit with attendees.

(L-R) Roberto O. Ramirez from NaSPA sponsor/advertiser Edwards Disaster Recovery Directory with Sharon and Leo Wrobel

For information on how to exhibit at next year's NDR show or for info on **how NaSPA members can attend at no charge**, email customercare@naspa.com.

Part 2 of 2—Performance Optimization Considerations in the Virtual Infrastructure Environment

Adapted from VMware ESX Essentials in the Virtual Data Center by noted authors David Marshall, Stephen S. Beaver, and Jason W. McCarty

In the last issue of Technical Support we discussed (1) How VMware ESX is one of the hottest server virtualization platforms on the market, (2) How it typically does a good job of providing high performance straight out of the box; but, (3) How even the best products still need performance tweaks on occasion. We continue therefore with our "tweaking tips" in this article, beginning with:

Host Server Storage Performance

When searching for a performance bottleneck in your virtualization system, the CPU and the RAM are usually the first two suspects. But by increasing the density and usage of a physical machine with numerous virtual machines, the amount of disk I/O activity greatly increases across that server's disk subsystem. And certain workloads are also very sensitive to the latency of I/O operations. Because of that, the host server's storage system becomes the third bottleneck that needs to be addressed in a virtual world.

Storage performance issues are usually the result of a configuration problem. With ESX, you have the choice of using either local or remote storage. Either selection comes with its own set of configuration choices to be made that can affect performance. And there are also a number of other dependencies such as workload, RAID level, cache size, stripe size, drive speed and more.

Without sounding too obvious, using the fastest and highest performing disks and controllers will greatly improve the performance of your virtual environment. If you have decided to go the route of local or direct-attached disk storage in your environment, you should go with 15K RPM disks to improve the I/O performance across all of your virtual machines on that host server.

In addition to the speed of the disk drive, you should also consider the type of disk drive. While SATA drives are now possible to use with the latest version of ESX, you are better off for performance reasons spending the money and selecting either Ultra320 SCSI drives, or even better if your system supports it, SAS disks.

Many disk controllers can support multiple channels on the card, and by splitting your disks across multiple channels, you can achieve a performance improvement. For example, if you have six SCSI disks and a two-channel controller, your hardware might allow you to place three disks on each channel and configure them in a six-disk RAID-5 array. This would allow you to effectively split your I/O across two channels.

You might also be able to install multiple disk controllers and additional disks within your host server. This would allow you to split up your file system and strategically place your virtual machines according to I/O needs. In other words, if you have ten virtual machines on a host server and two of those virtual machines are running disk I/O intensive applications, you have a choice on how to configure your environment for the best possible performance. You can either locate the eight low intensive virtual machines on one controller with the two high intensive virtual machines on the other controller, or you can configure five virtual machines on each controller where only one disk I/O intensive virtual machine is allocated to each file system.

TIP
Make sure I/O traffic isn't queuing up in the VMkernel. To verify, monitor the number of queued commands with esxtop over a period of time.

If your organization is lucky enough to afford one, it is recommended to use SAN storage rather than using locally attached disks so that you can achieve much better performance. SAN storage technology allows VMware ESX to shine with its added capabilities. Using a SAN will offload I/O operations from the ESX host server which leaves more resources available to the virtual machines. To further optimize performance, spread the I/O loads across multiple 4Gbps Fibre Channel SAN HBAs where possible. And make sure that heavily used virtual machines aren't accessing the same VMFS volume at the same time. Spread them across multiple VMFS volumes so that disk performance won't suffer.

VI3 has added new options for you to take advantage of in your remote storage. In addition to Fibre Channel SAN, you can now use iSCSI and NFS to take advantage of cheaper storage solutions using your existing IP networking technology. For iSCSI and NFS, it is important to make sure that your configuration and design does not result in an oversubscribed Ethernet link. The TCP session control will ensure that there is recovery from packet loss, but frequently recovering from dropped network packets will result in huge performance problems. Virtual machines with intensive I/O applications should not share Ethernet links to a storage device. And they will perform even better if they have multiple connection paths to its storage.

When your virtual machines share access to the ESX host server's I/O subsystem, use the I/O share allocation for each virtual machine to adjust the amount of I/O resources that the virtual machine is given. For virtual machines running applications that aren't very I/O intensive, you can set their resource shares to something low, like 500. And for the more resource intensive virtual machines that require more priority to I/O resources, set their shares to something higher, like 2000.

Host Server Network Performance

Network utilization can also present bottleneck issues in your environment much like CPU, memory and storage. But in most virtualized environments, you will find that the network is probably the least likely culprit of performance problems. However with that said, the host server still needs to be supplied with the appropriate amount of network bandwidth and network resources so that the virtual machines don't add any significant amount of network latency into the equation.

If you haven't already done so, upgrade your network environment with Gigabit Ethernet. With 10 GigE waiting to take over, Gigabit Ethernet network adapters and switches should be affordable. Using Gigabit network adapters allows more virtual machines to share each physical network adapter and it greatly improves the amount of network bandwidth made available to network intensive virtual machines.

When configuring your physical network adapters, the speed and duplex settings on each card must match the speed and duplex setting used on the switch port to which it is connected. The VMkernel network device drivers start with a default speed and duplex setting of auto-negotiate. The auto-negotiate setting is fine and should work correctly with network switches that are set to auto-negotiate. This is the default and preferred setting for gigabit connections. When using 100Mbit Fast Ethernet adapters, you should set the network adapter and the switch port speed and duplex settings to match at 100/Full. If you have conflicting settings between the network adapter and the switch, it can not only cause a performance problem but in some cases a connectivity issue as well.

You can also increase the available network bandwidth and increase network fault tolerance by teaming multiple Gigabit network adapters into a bond. This will also simplify the number of virtual switches being mapped to physical network adapters as well. You can also use separate physical network adapter and vSwitches to avoid network contention between the service console, the VMkernel and the virtual machines. It can also be used to segment network I/O intensive virtual machines from one another.

TIP

If you encounter bandwidth issues, check to make sure the NIC auto-negotiated properly. If not, change the speed and duplex settings manually by hard coding them to match. Do so at the switch or the VMkernel networking device using the VI Client.

You might want to leverage the new VMware ESX 3.5 networking enhancements that have been integrated into its networking code. Jumbo Frames are now supported. Supporting Ethernet frames up to 9000 bytes (as opposed to standard Ethernet frames supporting a Maximum Transfer Unit of 1500 bytes) allows guest operating systems using Jumbo Frames to require fewer packets to transfer large amounts of data. Not only can they achieve a higher throughput, but they also use less CPU than a standard Ethernet frame.

Another new feature in 3.5 is the support for TCP Segmentation Offload (TSO). TSO is widely used and supported by today's network cards. It allows for the expensive task of segmenting large TCP packets of up to 64KB to be offloaded from the CPU to the NIC hardware. ESX 3.5 utilizes this concept to provide virtual NICs with TSO support even when the underlying hardware doesn't have the special TSO capabilities. Because the guest operating system can now send packets that are larger than the MTU to the ESX server, processing overheads on the transmit path are reduced. TSO improves performance for TCP data coming from a virtual machine and for network traffic that is sent out from the server.

Configuration of the Virtual Machine

When you create virtual machines for your environment, accepting default responses aren't necessarily the right choice. Will it work? Sure. Will it be optimized for performance? Perhaps not. There are things that you should be aware of and things that you can do with the configuration of your virtual machines to squeeze out additional performance.

Remove Unneeded Virtual Hardware

One of the nice things about working with a virtual machine is the ease with which you can add or remove hardware components. You don't need any tools, and you don't need to open the hood of the server. A quick and easy way to gain a small amount of performance back in your guest environment is to disconnect or remove any unused or unnecessary devices. If your virtual machine doesn't need a CD-ROM drive, Floppy drive, Network Adapter or COM and LPT ports, get rid of them. When you do need them, it is extremely easy to either enable them or add them back. This can free up IRQ resources and it eliminates IRQ sharing conflicts.

Power off Idle Virtual Machines

Because virtual machines can so quickly and easily be created, a common problem called virtual machine sprawl has erupted. If it isn't managed properly with processes and controls, unmanaged and forgotten virtual machines can spring up and consume data center resources. It is important to identify virtual machines that are no longer necessary or no longer being used, because these virtual machines should be powered off or suspended to keep them from wasting valuable resources.

Even when a virtual machine is idle, a virtual machine is idle, it still uses memory and CPU that other virtual machines could be leveraging for their own performance.

Virtual Machine Processors

When a virtual machine's performance seems sluggish, many people automatically assume that it just needs more processing power. And since VMware has upped the ante on the number of virtual processors that can be assigned to any given virtual machine, why not just add more virtual CPUs to the virtual machine to increase performance? Unfortunately, that doesn't always work. In fact, in many cases, doing so may adversely affect the entire environment and further degrade performance.

Adding additional processors to a virtual machine by using VMware's Virtual SMP is not always going to solve the performance problem. Why? First off, your performance problem may not even be caused by a lack of processing power. Your bottleneck could be memory, disk I/O or network I/O. If the guest operating system or applications aren't being starved for CPU resources, adding more CPUs may only negatively affect the rest of the environment on that host server.

And second, not all operating systems and applications are able to take advantage of having multiple CPUs. If your application is single threaded and will not benefit from the additional VCPUs, you shouldn't use Virtual SMP. In this case, adding Virtual SMP may consume valuable physical processor resources without actually offering the virtual machine any added performance benefits, and as a result, take away resources from other virtual machines on the same physical host server. Virtual SMP should therefore be used sparingly and only after determining whether or not an application is indeed multi-threaded. Running a virtual machine with a single virtual CPU, in many cases, will outperform the same virtual machine with Virtual SMP turned on. This determination should be treated on a case by case scenario, and you should test each virtual machine environment to see if Virtual SMP will help or not.

Make sure to configure your virtual machines with the correct HAL or kernel. Single processor virtual machines should be configured with a UP HAL or kernel while a multi-processor virtual machine should be configured with an SMP HAL or kernel. SMP versions are required in order to fully utilize multiprocessor systems but can also be used on single-processor systems. Because of their synchronization code, SMP operating systems used on single-processor systems will operate slower than a UP operating system.

Another way to affect virtual machine performance is by altering the virtual machine's minimum and maximum CPU resource allocation percentages. If you want to avoid CPU starvation for a virtual machine, set its minimum percentage to something other than zero. Conversely, to keep low priority virtual machines from consuming too many CPU cycles, you can set its maximum percentage to something lower, like 50%, thereby effectively throttling down that virtual machine and allowing other more pertinent virtual machines to make use of those valuable CPU cycles.

In addition to setting thresholds with MIN and MAX settings, you also have the ability to control which physical processor or processors that each virtual machine can use. This control is called processor affinity. The default setting is to use no affinity, and this is usually the best choice for most situations. You should really only set a virtual machine's CPU affinity when absolutely necessary. As an example, if you have a very resource intensive virtual machine running on a host server, you might want to set its CPU affinity to isolate that virtual machine and to protect its performance. Doing so will also protect the performance of all other virtual machines running on that same host server by changing each of their affinity settings to a different processor from that of the resource intensive virtual machine.

Virtual Machine Memory

If you are trying to optimize the performance of your virtual machine and find that the performance degradation isn't being caused by the processor, examine the virtual machine's use of memory. If the guest operating system is paging or swapping memory too much, performance will suffer since writing to disk is much slower than writing to memory. To identify if your virtual machine is paging, use vmstat from the command line on Linux or the Performance tool found under Administrative Tools on Windows to check the value for pages/second. If the number is high, such as 1000 pages/second, increase the amount of memory assigned to the virtual machine to eliminate excessive paging.

Only allocate as much memory to the virtual machine that is needed to allow enough memory to hold the working set of applications being asked to operate in the guest environment. Some amount of testing is needed here since giving a virtual machine too much memory will reach diminishing returns and simply be wasteful. The wasted memory could have been used as an additional resource elsewhere on the host server.

If possible, configure Linux virtual machines to use less than 896MB of memory. Linux uses different techniques to map memory in the kernel if the amount of physical memory is greater than 896MB. Every physical page of memory up to 896MB is mapped directly into the kernel space. This memory section is faster and more efficient and can keep your guests running optimally. Any amount of memory over that is no longer permanently mapped but is instead temporarily mapped. These techniques add additional overhead on the virtual machine monitor and can lower performance.

NOTE: It usually depends on whether you are talking to the VMware ESX administrator or the virtual machine end–user as to what the "proper" configuration amount of memory is for each virtual machine. The amount of memory a virtual machine actually needs and the amount of memory the end-user wants normally doesn't match up. Finding a proper balance is the key to a successful infrastructure implementation.

Virtual Machine Networking

If you are trying to create a networking connection between virtual machines that live on the same host server, connect the virtual machines to the same vSwitch. While it isn't mandatory that these virtual machines be connected to the same vSwitch, doing so will keep the networking traffic from going out and across the wire. At the same time, it also keeps CPU and network overhead down and increases the performance of the network communications that take place between those virtual machines.

When you create a new virtual machine, the default network adapter that is emulated on a 32-bit guest is the AMD PCnet32 device and it is configured with VMware's vlance driver. For performance reasons, in a GigE environment, you should change the emulated network adapter to use either VMware's vmxnet driver or e1000. The vmxnet driver passes through network traffic from the virtual machine to the physical network adapter with minimal overhead. And the latest version of VMware ESX provides a new version of the vmxnet virtual device called Enhanced VMXNET. It includes several new net-

working I/O performance improvements such as support for TCP/IP Segmentation Offload (TSO) and Jumbo Frames. It also includes 64-bit guest operating system support. The vmxnet driver comes with the VMware Tools.

Configuration of the Guest Operating System

As mentioned earlier in this series, if you are trying to increase the performance of your virtualization environment, why wouldn't you attempt to tune your guest operating system? Just because this is a virtual environment, don't think that many of your normal operating system tweaks used in your physical servers isn't applicable here. They very well could be the case. There are numerous books, Web sites and other informational sources available to you that help with common operating system tweaks. This article cannot possibly cover them all. Instead, it will discuss some of the tuning options available that are specific to a virtual machine environment.

Updating VMware Tools

As simple as this may seem, it is important to remember to keep your VMware Tools up-to-date inside your guest operating system. You should always install the latest version of the tools in your virtual machine, and you should ensure that it gets updated after each VMware ESX upgrade or patch. There are times when update patches include fixes for components of VMware Tools, which would make these updates very important to your virtual machines. If you migrate or convert virtual machines from an earlier version of ESX or from another VMware virtualization platform, make sure to remove the old VMware Tools and then install the latest version for your platform. There are differences in the editions of VMware Tools from one VMware platform to the next. If you migrate or convert a virtual machine from another vendor's virtualization platform, make sure to remove that vendor's tools and replace them with VMware's.

By installing the VMware Tools in your guest operating system, you also get these performance enhancement capabilities:

VMXNET driver—As discussed earlier, VMware Tools provides a new updated and enhanced high speed networking driver to help improve networking performance.

TIP
Make sure to use supported guest operating systems. VMware Tools may not be available for an unsupported guest.

Improved video graphics driver—VMware Tools installs an improved graphics driver which gives you better mouse, keyboard and screen performance. In a Windows virtual machine, enable the hardware acceleration feature under the advanced graphics display settings on the Troubleshooting tab. This will help smooth out the mouse when remoting into the virtual machine.

BusLogic driver—Installing VMware Tools updates the BusLogic SCSI driver within the guest operating system to the VMware supplied driver. This VMware version of the driver has certain optimizations that the guest operating system supplied driver does not.

Memory balloon driver—The balloon driver is part of VMware Tools and is used for memory reclamation on ESX. It helps to minimize ESX swapping by better managing guest memory. Memory ballooning will not work without Tools installed.

Idler program—An idler program is added for Netware guests to help de-schedule these guests when they go idle.

Timer sponge—An experimental timer sponge has been added to give a correct accounting of time within the guests.

Time sync—When VMware Tools are installed, a time synchronization feature is added to improve the time keeping function within the guests.

Microsoft Windows® Guest Operating System Performance

Many normal tweaks and tuning techniques used to increase the performance of a Microsoft Windows environment on a physical server can translate over to a virtual machine. And just like a physical server, not every virtual machine is alike—so the tuning method used on one machine may not work the same on another.

The latest version of VMware ESX now supports 64-bit Windows guest operating systems as well as 64-bit applications. These 64-bit versions will typically have better performance than their corresponding 32-bit counterparts.

Defragmenting the contents of a virtual machine's hard disk can have a positive effect on its disk I/O performance. Using a third-party tool like Diskeeper can keep your disk structure well organized. Be sure not to schedule the defrag task during critical or normal business hours or else you might cause more performance problems than you fix.

You can free up resources within a virtual machine by stopping and disabling unnecessary services and background tasks. Make sure however that you don't disable a service that is needed by one of your applications. The following list contains a subset of common Windows services that can usually be stopped and disabled in a Windows virtual machine. This list should only be used as a suggestion until you know if a particular service is needed or not.

- Alerter
- Clip Book
- Computer Browser
- DHCP Client (Unless using DHCP IP addresses)
- Fast User Switching Compatibility (Windows XP)
- IMAPI CD-Burning COM Service (Windows XP)

- Indexing Service (Unless needed)
- Internet Connection Firewall (ICF) / Internet Connection Sharing
- IPSEC Services
- Messenger
- Network DDE
- Network DDE DSDM
- Network Location Awareness (NLA)
- Print Spooler (May be needed in some cases)
- Remote Desktop Help Session Manager
- Remote Registry (May be needed)
- Routing and Remote Access (May be needed)
- Smart Card
- SSDP Discovery Service
- System Restore Service (Windows XP)
- Telnet (May be needed)
- Themes
- Uninterruptible Power Supply
- Windows Audio (Windows XP)
- Windows Image Acquisition (WIA)
- Windows Time (May be needed)
- Wireless Zero Configuration

If your guest operating system is Windows XP, you may want to disable the System Restore feature. Doing so will free up system resources like CPU, disk space and I/O. Instead, use the more powerful VMware snapshot feature.

Keep your Windows virtual machine as lean as possible. Uninstall any of the Windows components that aren't going to be used. Doing so will reduce the amount of memory and CPU consumed within the operating system, and will offer more resources to the applications installed in the guest. Power features such as hibernation and hardware power management (turning off hard drives, monitors, etc) really don't have much meaning in a virtual machine or add any value. Likewise, screen savers, visual effects and animations consume additional CPU resources unnecessarily and should be disabled. In most cases, you probably don't need desktop wallpaper on your virtual machine either. One exception to that might be to use something like Microsoft's Sysinternal BgInfo tool which can be very useful by offering system metrics and information as the desktop wallpaper.

Linux Guest Operating Systems Performance

Many normal tweaks and tuning techniques that are used to increase the performance of a Linux environment on a physical server can translate over to a virtual machine. And just like a physical server, not every virtual machine is alike—so the tuning method used on one machine may not work the same on another.

TIP

The time synchronization feature is not turned on automatically. Be sure to enable time sync between the ESX host server and the virtual machine from within VMware Tools. When using VMware Tools to do the time synchronization, you should not use another form of time sync within the guest operating system. Do not enable VMware Tools time sync for Windows virtual machines that are members of an Active Directory because they should time sync with a domain controller instead.

The latest version of VMware ESX now supports 64-bit guest operating systems as well as 64-bit applications. These 64-bit versions will typically have better performance than their corresponding 32-bit counterparts.

The guest operating system timer rate found in the Linux kernel can impact the performance of a Linux virtual machine. Linux operating systems keep time by counting its timer interrupts. Unpatched 2.4 and earlier Linux kernels request clock interrupts at 100Hz or 100 interrupts per second. By default, some 2.6 Linux kernels request interrupts at 1000Hz while others do so at 250Hz. The overhead involved with delivering so many virtual clock interrupts can unnecessarily stress the virtual hardware and negatively impact performance. The only way to change the behavior is to reduce the number of ticks per second and then recompile the kernel. Where possible, it is best to try and use a Linux kernel that uses lower timer rates to avoid this problem.

If your virtual machine is running a server-class Linux guest operating system and X server is not required for the environment, disable it. Unnecessarily running X server and screen savers in your guest environment will consume extra resources that can negatively affect performance of the virtual machine or its host server. When you do need a graphical desktop, make sure you select a light-weight window manager.

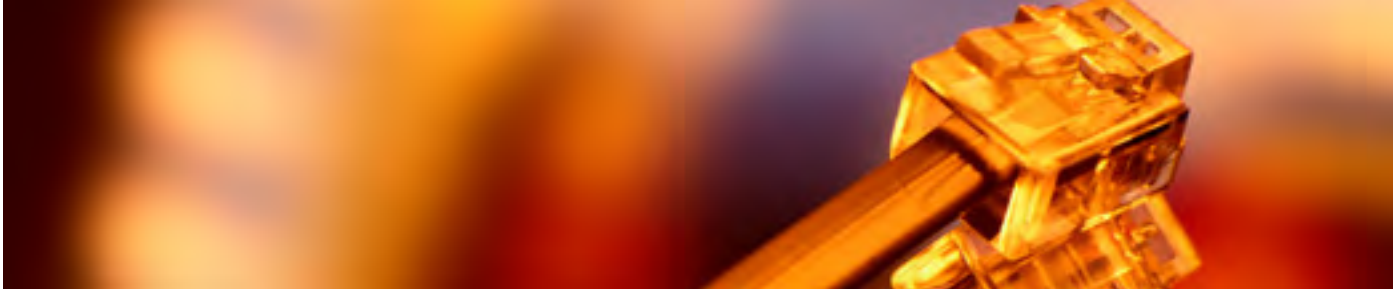

When optimizing your Linux virtual machines, make sure that you disable or remove any unnecessary daemons, services and background tasks. At the same time, be sure to remove any unneeded packages as well. This will free up processor and memory resources that can be redirected to your applications.

Backups and Anti-Virus Configuration

Backup and anti-virus configurations will probably vary based on the use of your virtual environment. In other words, how you configure backups and anti-virus to improve performance in your environment will depend on if your virtual infrastructure is a production environment or a non-production environment such as test and development. A non-production environment may not need to have anti-virus software installed (unless you are testing the anti-virus solution or the effects of anti-virus solutions). However, if this is a production environment running a Windows guest operating system, you will probably be required to have an anti-virus package installed. In that case, it should be configured in the most optimal way in order to help maximize the performance of the virtual machine.

One of the first things you can do is to setup scheduled virus scans to take place during off-peak hours so that the applications in the virtual machine do not compete for resources with the anti-virus solution. Unless your organization specifies otherwise, in most cases, a single, daily virus scan should be adequate. Additionally, most servers typically do not need real-time virus scanning turned on since this can greatly impact server performance. Disable the real-time virus scanning features when possible for production virtual machines, especially those used as database or web servers. Also, most anti-virus solutions have the ability to exclude certain files, file types, and directories from being scanned. Some application file types, such as database data files, do not need to be scanned, and in fact doing so could destroy performance. You should also configure the anti-virus solution to exclude mission critical application files that are not high risk for virus infection. Swap files are also good candidates for exclusion.

If you are running a backup agent in your virtual machines, you should also schedule the backups to take place during off-peak hours, and equally important, not during a scheduled virus scan. Again, just as with anti-virus software, scheduling backups to take place during off-peak hours will help alleviate performance degradation from applications competing for resources with the backup solution. If your organization is using VMware VI3, it's also a really good idea to make use of VMware's Consolidated Backup feature or use a third-party application that leverages the technology. Doing so will help eliminate a lot of performance overhead within the virtual machines and on your network because the technology will offload much of the overhead onto your SAN and off of your host server.

Summary

If you decide that you don't want to perform any other optimization or tuning suggestion found in these articles, at least consider upgrading your environment to the latest version of VMware ESX. Doing so will offer you immediate environmental optimization and give you additional tools with which to better operate and control your virtual infrastructure. Upgrading can make it even easier for organizations to virtualize their most resource intensive and demanding applications. Otherwise, each of these best practice performance enhancements can be viewed on a case by case basis as needed.

We hope you enjoyed this series and that it was helpful to you. If you really liked it, why not consider buying the rest of our book with your NaSPA member discount? Not only will it help you professionally, it will help your NaSPA as well since a portion of every sale goes to NaSPA. If the concept works, NaSPA will continue to bring you, the member, quality content like this for free, and meaningful discounts from great authors and publishers in case you need to get into more detail. Until next time, thanks for reading.

With Group Savings Plus®, NASPA members can get more from their auto and home insurance.

Responsibility. What's your policy?™

Savings of up to $327.96 or more a year on auto insurance*
with a special group discount and other discounts**

12-month Rate Guarantee
unlike the six-month policies that some other insurers offer

Help when you need it
with 24/7 Emergency Roadside Assistance and 24-hour claims service

Additional coverages for added security
including Umbrella Liability policies, Accident Forgiveness† and Home Insurance with optional Identity Fraud Expense Coverage

Chances are, Liberty Mutual's Group Savings Plus program may be able to offer you more savings and benefits than your current auto and home insurance provider.

AUTO

HOME

Get more. Save more.
Find out just how much more today.

- **Call** 1-800-524-9400 and mention client #101282
  Monday-Friday, 7 a.m. to midnight; Saturday, 7 a.m. to 9:30 p.m.; Sunday, 10 a.m. to 8 p.m. (ET)
- **Go to** www.libertymutual.com/naspa
- **Or visit** a Liberty Mutual office near you

This organization receives financial support for allowing Liberty Mutual to offer this auto and home insurance program.

*Figure based on a February 2008 sample of auto policyholder savings when comparing their former premium with those of Liberty Mutual's group auto and home program. Individual premiums and savings will vary. **Discounts and credits are available where state laws and regulations allow, and may vary by state. To the extent permitted by law, applicants are individually underwritten; not all applicants may qualify. †Accident Forgiveness coverage subject to terms and conditions of Liberty Mutual's underwriting guidelines and is not available in all states. Coverage provided and underwritten by Liberty Mutual Insurance Company and its affiliates, 175 Berkeley Street, Boston, MA. A consumer report from a consumer reporting agency and/or a motor vehicle report, on all drivers listed on your policy, may be obtained where state laws and regulations allow. Please consult a Liberty Mutual specialist for specific details.

Situational Awareness: New Buzz Word for Contingency Planners?

By Karie Wohlgemuth, NC4

With all of the negative news about the economy and words like "downsizing" and "furlough days" becoming the new buzz words, companies are looking for ways to cut costs but still protect their business from unexpected events. More and more that effort means utilizing technology to supplement a smaller work force, while still managing corporate risk. The first and most important part of significant event management is in knowing what happened and where. Set in this context, it might be a good idea to consider what the level of Situational Awareness is in your organization. Have you checked lately?

It Starts With Knowing What's Going On

BOOM! What was THAT?

Any disaster response must start with knowing *to a certainty* precisely what is going on. It is only through effective collection and dissemination of accurate information that an effective disaster response can be mounted. In the past, NaSPA has discussed the "dissemination" part of the equation in some fine articles about outbound notification systems. But what about the "what happened?" or more specifically, "what's happening?" part of the equation? If a disaster strikes, how will you know *exactly* what has happened, when half facts and misinformation may abound? How does one afford to even propose such a capability these days? The answer is that as in most of the other aspects of our business, technology is being employed. Today's technology not only warns and alerts organizations to disasters, a fire alarm can do that on an elemental level. Today technology can "filter out the noise" by sifting through thousands of alerts and only telling you about the important ones. The net result is the ability to deliver *actionable* information to your organization based on the situations and scenarios that are most meaningful to it.

Are Companies *Really* Doing This Today?

Absolutely, and the bigger they are the more they do it. After 9/11, it became obvious that companies needed better ways to see the events unfolding around them, and to communicate with employees, stakeholders and partners. Yet, it is virtually impossible to monitor all of the information sources that are available on a 24/7/365 basis. Even if a company could monitor all of these sources, it would waste resources sifting through the volumes of data only to mine the small nuggets of information that may actually be applicable. A car alarm in the parking lot may not be important. An explosion in the basement probably is. A train derailment might not be important, unless it's next to your place of business and spills a load of ammonia.

Consider the incident on April 9, 2009 when someone sabotaged an AT&T fiber optic cable and blacked out a big chunk of California. The

deliberately severed fiber optic cable caused millions in Silicon Valley and San Francisco to go without broadband, phone, and wireless service for most of that day.

In this case, *Twitter* was actually used by communications giant AT&T to "reached out and update" its customers about the sabotage! AT&T began "tweeting" about the fiber sabotage in California around 7 a.m. PDT on April 9, 2009 with the first message saying:

"CA customers: We are aware of a cable cut situation impacting services in Santa Clara and San Jose areas."

AT&T went on later to add:

"AT&T offering $100,000 reward for info leading to arrest/conviction of those responsible for CA vandalism. Call 408-947-STOP."

When you think about it, however, wouldn't one assume there were better ways to organize information and communicate in a time like this than with a social networking tool? As it turns out, there is and some companies have been developing the tools for some time.

The Industry's Response to Situational Awareness

In a research expedition, members of the Candle Corporation (now part of IBM) went out to the marketplace to see if technology could answer the call for being the intelligence and reporting arm for incidents happening externally to an organization, that might impact its people and operations. There was not only a need for such a technology, but a gap in the marketplace in fulfilling this need. As a result, NC4, the National Center for Crisis and Continuity Coordination, was formed in 2002. The focus of NC4 was not only on creating awareness in the marketplace, but also on readiness for the things happening around you. NC4 created the phrase, "Situational Readiness" to apply to its offering.

To be effective, NC4 would have to monitor thousands of the media sources, emergency service scanners, dispatch systems and social media sites around the world while filtering the information to determine relevance to a customer. The vision was to be the primary information source for relevant events that had a direct impact on customers. This vision translated into the External Situational Awareness service, known as ESA. NC4 identified the types of information that corporations and government organizations would be interested in which now include: advisories, aviation, fire, geospatial, hazmat, health, infrastructure, meteorological, security, structural, terrorism and transportation.

It is just as critical to understand an event's proximity to a specific location as it is to understand the severity of the incident. I may not be interested in a traffic accident 50 miles away, but I will be interested if it's outside my building and affecting my employee's ability to leave work or my distributor's ability to deliver supplies to my store. Similarly I would want to know if severe weather is headed in my direction or if a terrorist activity affects the hotel where my employees are staying. NC4 customers therefore pinpoint locations on a map and then decide the incident type, proximity and severity that are applicable to their specific organization. This allows them to then profile what they want to know about, as well as the devices that they want used for notification when a given event triggers that profile. The key here is that relevant alert information is filtered and distributed based on the criteria they define. The information can be received over e-mail or any SMS device which means customers can see the incident summary on a cell phone, pager or blackberry.

Does it work? According to Ray Ferrara of Ferguson Enterprises, "Everyone asks about return on investment (ROI) after installing a technology solution. It is very difficult to quantify the dollars saved from preventing a terrible event, making decisions five minutes faster or knowing that your employees are safe. The peace of mind ESA provides is priceless." Similar stories can be heard from organizations that use this technology to provide them with critical information to help them maintain business continuity. NC4 customers include major financial firms, aerospace and defense contractors, Government agencies, software and technology vendors and pharmaceutical companies.

Yet, the basic service that NC4 set out to create 7 years ago is a much different product than what ESA is today. NC4 still maintains the top-notch incident monitoring and alert notification that made it an almost instant success. NC4 has partnered with local police, corporate security and emergency services to receive direct information feeds as well as incident confirmations from these sources. This growing list of direct sources, in addition to the extensive technology developed specifically for ESA, gives NC4 advantages that traditional alarm or monitoring services do not have. NC4's technology adds automated sources as well as parses the data into bytes that are monitored by one of the analysts in the NC4 Incident Monitoring Centers (NIMC.) The service has both a domestic and global offering.

In following their mission of using technology as a force multiplier, NC4 has partnered with TranSecur to expand the offering to include

Want to Know More?

Other recent events, news articles and press releases can be found at: www.nc4.us/newsevents.php. Information on ESA including Case Studies: www.nc4.us/ESA.php Incident Highlighter: www.nc4.us/flashmap.php (Shows a preview of the live incidents being monitored by NC4 Incident Monitoring Centers) Also see the NC4 ad in this section. We thank Karie and NC4 for their thoughtful contribution to Technical Support Magazine*, just in time for National Preparedness Week.*

country risk analysis and will introduce its travel tracking component later this year. Also, NC4 has recently completed its integration with Everbridge to allow incident data to be broadcast to tens or thousands through Everbridge's extensive notification distribution technology. The ESA service now has a sister component, ISA which is the Internal Situational Awareness feature that allows companies to report, track and distribute information regarding their own internal incidents in addition to the external incidents reported by NC4. All of these features can also be served up through a mobile device.

NC4 has added to its technology suite the E Team emergency response platform, which allows emergency services to prepare, respond and recover from natural or man-made disasters. In addition to its domestic use by most major U.S. cities, E Team has also expanded into several other countries. Further, NC4 provides a secure collaboration portal product through its ESP Software as a Service package. This SaaS product is in use by many Federal agencies for secure collaboration and communication. NC4 is also presently evaluating other sources of useful data, such as the Pacific Disaster Center and others.

Summary

In closing, the use of technology can indeed make a significant impact on a business in time of distress if it is timely, relevant, filtered and well-utilized.

As is virtually always the case, technology should be employed where possible to do the heavy lifting. Besides, it's cheaper than just using people power.

Now is the perfect time to look at technologies with multiple uses that can bring in the benefits of risk mitigation and safety for employees and operations while leaving your resources to focus on other aspects of the business.

Karie Wohlgemuth is the Director of Marketing for El Segundo CA based NC4, which specializes in all facets of situational awareness for major corporations, and proud NaSPA sponsor. The company's web site can be found at www.nc4.us.

Security Professionals: Enhance Your Career Prospects by Developing a "Personal Brand"

By Jeff Snyder, President, SecurityRecruiter.com

Can a security professional create a personal brand? "Why SURE!" claims Jeff Snyder and we agree. And while this article from NaSPA's resident Security Recruiting expert deals prinicpally with people seeking employment or career enhancement in the Security profession, the tips are solid and apply to almost anyone in any profession. Read and heed as Jeff explains how to build your own "personal brand" in order to enhance your chances of gainful employment, in whatever profession you are seeking.

Personal Branding

Any security professional can build a personal brand and it is important to do so if you want to line yourself up with the best security jobs that emerge in the future.

Wikipedia suggests that **Personal Branding** can be defined this way: "Personal branding is the process whereby people and their careers are marked as brands."

With this definition in mind, here are some ideas for building your personal brand.

1 Your resume is a significant part of your personal branding package

While it isn't the only piece of your personal branding package, your resume needs to be carefully aligned with your skills, abilities and accomplishments. Your resume needs to clearly demonstrate how you've added value to previous employers. Anybody can write a resume that describes what they were hired to do. When a resume clearly demonstrates the value you as the owner of the resume has provided for past employers, future employers can then start to environ how hiring you might bring value to their organization. Your future employers are trying to get more done with less. Build your resume so it clearly shows value and is not just a snapshot of the job description you signed up for with past employers.

2 Build a blog focused on your area of expertise

If you're an expert in Identity Management, there is plenty to write about if this topic defines your expertise. Use your blog to display your subject matter expertise in the subject. Interact with those who might want to tap into your expertise by way of acknowledging comments left on your blog. Write in your blog with the intention of sharing your knowledge for the benefit and education of your readers. Sooner or later, your blog may be found by someone who not only appreciates your expertise but they may also have a job for you.

3 Write articles and get them published

My first full article was written for a major security industry magazine. Once this article was published, I shared it with other major industry publications. Not every publisher got excited about my first article but that didn't matter. Two major publishers appreciated my work and gave me an opportunity to write for them. Now I not only have a growing portfolio of articles but occasionally unsolicited invitations to write new articles come my way. Additionally, I built a Security Articles section on SecurityRecruiter.com where security professionals can get their articles published. Any security or business focused article that will bring value to the visitors to SecurityRecruiter.com will be considered. I'll even push your article out in my blog to get it more attention. (Note, NaSPA *Technical Support Magazine* is also always seeking new authors with meaningful things to say! E-mail Leo Wrobel, Executive Editor, president@naspa.com)

4 Speak in public

Sure, you might have your sights on a speaking role at a major conference. Start small by speaking at a local ISSA (Information Systems Security Association) or ISACA (Information Systems Audit and Control Association) meeting. Speak to a class of future security professionals in a university class setting. In the same way that you demonstrate subject matter expertise through a blog, demonstrate the same expertise by speaking in public. Public speaking is an outstanding way to enhance your own communication skills. Communication skills are critically important skills from the mid-manager to executive levels for security professionals.

5 Volunteer your time with a non-profit organization

Non-Profits, community organizations and churches have data and/or facilities that need to be secured. Maybe they can't pay you but you'll build up expertise and if you're lucky, a letter of recommendation. Once you have collected several recommendation letters, you can add these letters to your portfolio.

6 Build on-line network presence carefully and strategically

LinkedIn is the first place I suggest for building a professional business profile on-line. With over 40 million members, LinkedIn is still likely the fastest growing on-line business network in the world. Load your on-line profiles with the same kind of information you would use to build your resume. Make it easy for employers and recruiters to find you. Be conscientious about what you write on-line. Once you write and push your thoughts out to your audience, you can't take anything back in the future.

Jeff Snyder is the President of SecurityRecruiter.com, a search firm highly specialized in information security recruiting. Jeff's recruiting career started in 1990 in the general IT recruiting space. His first information security recruiting assignment landed on his desk in the 1995 - 1996 timeframe. SecurityRecruiter.com provides recruiting services, job placement services, professional resume writing services and is a gateway to various kinds of security certifications and training opportunities.

Smooth Moves

By Philip Jan Rothstein, FBCI

- Relocating a data center can be enough of a challenge, doing so with minimal business disruption demands good luck and great planning.
- Some of the most successful moves have been little more than activations of customized versions of the disaster recovery plan!

Moving a business? It can be an unsettling experience, to say the least. Moving a vital technology facility such as a data center or communications hub can be downright frightening.

I remember my father's advice a long time ago when I made the big move from renting an apartment to home ownership: He said to figure all your expenses, allow for problems or surprises, then add fifty percent to your budget — and effort. Any project manager moving a data center would do well to consider those words of experience.

A data center relocation, especially if construction is involved, is at best a twelve-month process. This is not to say it cannot be done faster; however, the direct costs along with resource commitments, and level of effort are likely to escalate almost exponentially with a shortened timeframe, or when other critical priorities divert essential resources.

Factors Foretell Fate

Immediately upon committing to a data center move, four variables can potentially make or break the entire project. The first and most critical action is choice of an effective project manager. Often, the data center manager is assigned by default. The only way he or she can tackle a move project is by handing off most, if not all, day-to-day responsibilities. Having guided about fifty data center moves, I have observed that failure was often a possibility when the organization refused to dedicate key management and staff full-time.

The second most critical action involves blocking out windows of risk or opportunity. Long lead-time factors which may impact project timing or the actual cutover -- the point during the move transition when operations are suspended at the old site until they resume at the new data center - must be identified and addressed immediately. On the other hand, some of these same factors, as expiring equipment leases, may offer timely opportunities for the relocation

Considerations include business cycles; real estate leases; construction schedules; equipment leases and depreciation schedules; equipment, software, network migration or upgrade plans; workload and capacity projections; MIS plans; and space constraints.

In most cases, a best-case cutover window of opportunity will be obvious from this analysis and will drive much of the project timing and resource levels. Unfortunately, relocations are often driven by non-negotiable factors.

Smart Moves

- **Allow sufficient time, typically 12-18 months**
- **Secure adequate funding, appropriate budgeting**
- **Commit dedicated staff**
- **Avoid or circumvent critical-path tasks**
- **Check early for "drop-dead" dates**
- **Treat a relocation like a disaster recovery**

The third critical success factor is tangible management commitment. Budgeting and resource allocation for a capital project as complex as a data center move must be identified and approved early, even though accurate budgeting tends to be difficult in the early stages. This must be recognized by top management and assurances obtained that funding and staffing will not be jeopardized once the commitment to relocate is made.

The fourth critical success factor involves taking advantage of concurrent opportunities. Simultaneous developments may include upgrading of aging equipment or equipment nearing lease expiration; technology enhancements which could dramatically reduce space and infrastructure requirements; or, staffing plans which might necessitate leasing additional office space. Such side benefits often can offset many of the costs of a data center move.

Declare A Non disaster

A typical data center build/move project may include more than a thousand discrete tasks and directly involve from ten to a hundred individuals: project manager; project secretary; user liaison; vendor liaison; software team; hardware team; finance; voice communications; data communications; insurance and risk management; physical security; information security; human resources; construction manager; general contractor; subcontractors; logistics specialist; customers;

numerous vendors; administrative team; architect; engineers; movers; salvage team; etc.

If it is not already obvious from that list, you will find a data center move process to have a lot in common with a data center disaster recovery. Some of the most successful data center relocations have been little more than activations of the operation's disaster recovery plan. As a bonus, this approach affords a valuable opportunity to exercise the recovery plan under realistic conditions. Exercises or dry runs, especially if they are based on a solid disaster recovery plan, are invaluable.

Moving Right Along

It's imperative that relocation activities interfere only minimally at the cutover. Project activities should be kept to an absolute minimum during this extremely vulnerable time period. In other words, any tasks which can be initiated and completed well before a 'freeze' window' can help to reduce distractions and risk.

Excessive downtime risk can be avoided, for example, by relocating or swapping equipment outside of the cutover timeframe. Ideally, a cutover should require very little physical movement although, of course, there will be functional as well as economic tradeoffs. A thorough inventory and "spring cleaning" can help to avoid nasty surprises as well as nuisances as the cutover looms.

Early and regular communication with every affected or involved party is crucial. There must be no surprises for data center customers, vendors, management, MIS staff or anybody else. This communication should begin on a monthly basis once the project is committed. By the week leading up to the cutover, daily communications are appropriate; from the day prior to the start of the cutover to the day after acceptance, hourly or even continual advisories are warranted. While I have not yet used Twitter during a relocation, it seems to me this could be a valuable tool. Progress reporting to management is also essential.

Tending To Timing

As the cutover approaches, a "freeze period" is an important mechanism to reduce risk and maintain sanity. Typically, for a period on the order of three weeks prior and two weeks after the cutover, business areas, MIS and others must recognize that any significant activity or change not only jeopardizes the data center move, but may in itself be at risk. Freeze period exceptions are to be expected, but must be closely managed; exception guidelines and procedures should be spelled out clearly in advance. Often, top management must be called on to intervene.

Exercises are as essential to a data center move as they are to a disaster recovery. Beginning at least a couple of months prior to the cutover, tabletop walkthroughs of the cutover should be held with the entire project team, including vendors and affected customers. Dress rehearsals will identify many potential glitches and smooth over rough spots. Again, the disaster recovery plan should be a guide – if the DR program isn't up to snuff, then it should be upgraded well before the cutover.

Timing of the actual cutover process is scheduled with military precision. Staffing schedules are mapped out by location, hour-by-hour. Backup and contingency staffing is on standby. Supplementary communications and transportation are in place. Tasks may be mapped out in quarter-hour increments. Decision and commitment (go/no-go) points, based on validation criteria, are explicitly defined.

Like a disaster recovery process, it cannot be assumed that all the ingredients will easily fall into place, that the other team is firing blanks and not live ammunition. The ammunition is going to be live and painful, if not fatal, if precautions for every contingency are not mapped out carefully in advance. These contingencies should have already been examined in dress rehearsals.

Moving physical equipment is typically a visible yet relatively minor aspect of the overall data center move project. In most cases, network connectivity and data storage transitions are more critical. As an example, climate-controlled vans and four-wheel drive tow trucks have been employed to ensure safe arrival of equipment and media; entire, parallel networks have been built to guarantee connectivity.

Making The Cut

The "right" strategy for one organization may be wrong for another. It is vital to focus on a strategy early in the planning process, although flexibility may be required as business, technical, construction or other considerations come into play.

One common move strategy, the "Flash Cut," involves shutting down the old center and moving the hardware, media, materials and so forth all at once. For smaller centers or low-risk operations which can afford a few days of service interruption, this approach may work well. Of course, this type of relocation can be riskier than most. If all goes well, expense should be low. If all does not go well expenses probably will not matter much.

At the opposite end of the spectrum, a data center move may involve minimal, if any, physical relocation. Redundant or replacement hardware and other components at the new site can be fully tested and operational before the cutover. This approach works especially well if a concurrent transition to replacement equipment is practical. From a disaster recovery perspective, this is pretty much equivalent to a "failover."

Key Stages of the Relocation Project

- Formation of a Project Team
- Initial Strategy Commitment
- Resource Allocation and Budget

- Developing the Relocation Workplan
- Committing to a Cutover Schedule
- Set a Freeze Period
- Early Handling of Non-Critical-Path Tasks
- Cutover and Acceptance
- Post-Cutover Phase

Short-term equipment rentals provide the opportunity to shake down the new facility and network with little risk. Rental equipment may be installed at the new facility prior to the cutover and used in production for a short period, or may serve as fallback while production equipment is moved.

In some cases, the rental equipment has even been installed at the old facility. Of course, its use for live operation during the move introduces an additional cutover step, from the rental back to the production equipment. In most cases, a blend of strategies is most effective. Some less critical components may be flash cut over a weekend, for example, and the most critical components are either replaced, or short-term rental equipment is employed.

How'd We Do?

I recall one data center cutover which went remarkably well until the CEO showed up at the door of the new facility, rather upset. It turned out that the one step skipped was notifying top management that we were done! In fact, the new data center had been up and running "live" for more than twelve hours and all data center customers were back to normal. The cutover had gone so smoothly that the CEO was convinced we had failed.

Like a disaster recovery process, a data center relocation project needs an explicit conclusion. A clean-up and salvage stage should not be ignored, as the former site may still require attention and there may be temporary equipment or other matters to resolve.

The shakedown period may require extra attention. A post-mortem review is often helpful. And, a party for the project team is definitely called for!

Philip Jan Rothstein, FBCI (pjr@rothstein.com), is president of Rothstein Associates Inc. (www.rothstein.com), a management consultancy founded in 1985 focused on business continuity, disaster recovery and crisis management. He is the author, editor and publisher of over 80 books, and Publisher of the Business Survival™ Blog as well as The Rothstein Catalog on Disaster Recovery, the industry's principal resource for books, training materials, videos and software. He was elected Fellow, The Business Continuity Institute in 1994 in recognition of his contributions to the industry.

Plan For Perils, Pitfalls

- **With such a complex undertaking, it is easy to be lulled into complacency and overlook a fatal flaw. Actual examples of serious problems experienced during past data center relocations include:**
- **A labor union jurisdictional dispute over inventorying and packing of computer tapes;**
- **Delayed freight elevator availability because of another building tenant, followed by the discovery that certain equipment would not fit in the elevator;**
- **A flu epidemic affecting most of key members of the move team;**
- **An irreplaceable computer component tipped over and damaged;**
- **Employee cars parked on the street overnight towed away;**
- **Hardware, software and communications not meshing together after a move, without a contingency plan to return to the former site;**
- **No certificate of occupancy for a new facility as of the committed move date, resulting from poor coordination with construction schedules;**
- **Resignation of the relocation project manager a week before the scheduled move, date without a qualified and up-to-speed backup in place;**
- **Significant configuration changes not accurately mapped into the move plan or new facility;**
- **Applications software changes "snuck in" three days before a move;**
- **Project team members immediately dispersed to other priorities right after the cutover and were unavailable when a cutover-related major network breakdown threatened the new data center.**

JOIN NaSPA TODAY
www.NaSPA.com

Please complete this form and fax it to (414) 768-8001

❑ **Yes, I want to become a NaSPA member!**

Name _____

Signature _____ Date _____
required

Title _____ Company _____

E-mail Address _____

Address _____

City _____ State/Province _____ Zip _____

Country _____ Phone (____) _____

Address shown is: ❑ Home ❑ Work

System software you work with: (check all that apply)

1. ❑ MVS-IBM 2. ❑ VM-IBM 3. ❑ VSE-IBM 4. ❑ Windows Workstation
5. ❑ Windows-Server 6. ❑ UNIX-All 7. ❑ UNIX-Linux 8. ❑ NetWare-Novell
A. ❑ AS/400(OS/400)-IBM B. ❑ Macintosh-Apple
C. ❑ VMware Z. ❑ Other _____
please specify

Payment Method:

❑ Check ❑ MasterCard ❑ Discover ❑ Visa ❑ American Express

Card # _____ Exp. Date _____

ADVERTISER DIRECTORY

SERVICE DIRECTORY

MEMBERSHIP INFORMATION
- Address Changes
- Dues Payment
- Membership Status/Applications
- User ID/Password Inquiries
- Publication Back Issues

Contact Carrie Banacin
Membership Manager
(414) 908-4945, Ext. 116
NaSPA_membership@NaSPA.com

PUBLISHING INFORMATION
- Article Submissions
- *NaSPA E-News*
- Letters to the Editor
- Copyright Inquiries

Contact Leo A. Wrobel
Editor-in-Chief
President@NaSPA.com

CUSTOMERCARE CENTER
Contact Bonnie Kuchinski
(414) 908-4945, Ext. 450
customercare@NaSPA.com

MEMBER SERVICES/CHAPTERS
- Chapter Start-up Kit
- NaSPA Direct
- Discount Programs

Contact Alex Llanas
Marketing Coordinator
(414) 908-4945, Ext. 123
a.llanas@NaSPA.com

ADVERTISING SALES
- Advertising in *Technical Support*™
- *NaSPA E-News* sales
- *Reprints*

Contact Don McMurray
Sales Director
(414) 908-4945, Ext. 111
dmcmurray@NaSPA.com

NaSPA INFORMATION
- Board of Directors
- NaSPA Focus and Direction
- Industry Issues and Trends

Contact Leo A. Wrobel
President
president@NaSPA.com

MEMBERSHIP CATEGORIES (CHOOSE ONE)

❑ **Subscriber**	• *Technical Support*™ and *Virtualize!*™ magazine • NaSPA luggage tag • *Rate per year: $40 U.S., $45 International*
❑ **Gold**	• *Technical Support*™ and *Virtualize!*™ magazine • Member ID Card & Certificate of Membership • Members only web access • All member discounts and services • One e-mail account • NaSPA luggage tag • *Rate per year: $45 U.S., $60 International*
❑ **Platinum**	• *Technical Support*™ and *Virtualize!*™ magazine • Member ID Card & Certificate of Membership • Members only web access • All member discounts and services • Two e-mail accounts • NaSPA luggage tag • *Rate per year: $95 U.S., $110 International*
❑ **Faculty**	• *Technical Support*™ and *Virtualize!*™ magazine • Member ID Card & Certificate of Membership • Members only web access • All member discounts and services • One e-mail account • NaSPA luggage tag • *Rate per year: $20 U.S., $20 International*
❑ **Retired**	• *Technical Support*™ and *Virtualize!*™ magazine • Member ID Card & Certificate of Membership • Members only web access • All member discounts and services • One e-mail account • NaSPA luggage tag • *Rate per year: $20 U.S., $25 International*
❑ **Student**	• *Technical Support*™ and *Virtualize!*™ magazine • Member ID Card & Certificate of Membership • Members only web access • All member discounts and services • One e-mail account • NaSPA luggage tag • *Rate per year: $20 U.S., $26 International*
❑ **Group**	• *Technical Support*™ and *Virtualize!*™ magazine • Member ID Card & Certificate of Membership • Members only web access • All member discounts and services • NaSPA luggage tag • One e-mail account per member • *Rate per year (for up to five members): $150 U.S., $150 International*
❑ **Life**	• *Technical Support*™ and *Virtualize!*™ magazine • Member ID Card & Certificate of Membership • Members only web access • All member discounts and services • Two e-mail accounts • NaSPA luggage tag • Choice of one NaSPA appreciation gift* • An individual Web site account with 20mb of space • *Rate: $995 U.S., $1,010 International*
❑ **Corporate**	• *Technical Support*™ and *Virtualize!*™ magazine • Five memberships • Member ID Card & Certificate of Membership • Members only web access • All member discounts and services • NaSPA luggage tag • Unlimited Job Postings for one year • Choice of one NaSPA appreciation gift* • *Rate per year: $1,850 U.S., $1,850 International*

*Gifts are available at the time of paid membership renewal only.