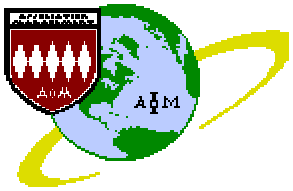

Journal of Information Technology Management

ISSN #1042-1319

A Publication of the Association of Management

APPLYING THE SCALE-FREE DEGREE DISTRIBUTION ALGORITHM TO ASSESS COMMUNICATION COMPLEXITY AND FAILURE POINTS IN DISASTER RECOVERY MODELS

OLIVIA F. LEE
ST. CLOUD STATE UNIVERSITY
oflee@stcloudstate.edu

DENNIS GUSTER
ST. CLOUD STATE UNIVERSITY
dcguster@stcloudstate.edu

MARK B. SCHMIDT
ST. CLOUD STATE UNIVERSITY
mbschmidt@stcloudstate.edu

BRANDON MCCANN
ST. CLOUD STATE UNIVERSITY
mcbr0303@stcloudstate.edu

ABSTRACT

In the current digital era, organizations are increasingly aware of the need for a successful disaster recovery (DR) plan. DR ensures the continuation and progression of business regardless of a single or multi-point failure by offering redundant systems and multi-point backups to help ensure a successful recovery. A carefully designed DR plan is crucial for recovering information, and a vital strategy for sustaining daily operations. Although past research has discussed many recovery sites options, the understanding of recovery site communication paths and their associated complexity is still limited. Using the scale-free degree distribution formula, the authors present a methodical discussion concerning the network characteristics of various disaster recovery options, and the investment required for technology infrastructure and personnel support within various models. The current study marks a pioneering effort in the DR field by applying the scale-free degree distribution formula to assess the network complexity index and overall model failure points. In addition, a modified hot site designed especially for small and medium size businesses is presented to leverage inexpensive commercial hardware by using standard PC components. Some important implications of this paper include devising a practical assessment tool for DR planning, recovery investment analysis through comparison of recovery site spending, and infrastructure requirements for various recovery options.

Keywords: disaster recovery, degree of distribution, recovery site options, network complexity, failure point.

INTRODUCTION

Interruption of normal business operations can occur for many reasons. Some possible causes of interruption to normal business operations are natural or man-made disasters. Indeed disasters can occur in many forms, such catastrophic events as hurricanes, tornados, floods, power outages, fire, vandalism and theft. In the realm of electronic data recovery, many options are available to protect data and ensure business continuity in the wake of the crippling damage a disaster may cause. Backup systems, such as tape drives, external removable hard drives and similar devices, are simply not effective, if the data is not protected offsite. Having a plan in place to transition control of business operations, for even a brief period when a disaster strikes, can save thousands of dollars of unnecessary expenses.

Disaster recovery (DR) ensures the continuation and progression of business regardless of single or multi-point failures by offering redundant systems and multi-point backups designed to provide successful recovery. Quick recovery is vital to avoiding the crippling effects associated with a work stoppage. The objective of a disaster recovery plan is to return the business to an operational status equivalent or better than before the disaster occurred. Thus, recovery efforts must ensure any failure can easily be reversed and critical information can be restored to the point where business may take place as usual. This is an important measure to make certain repeated disasters do not cause degradation of business performance as a whole. Prior to developing a DR plan and selecting recovery site options, it is necessary to conduct a business impact analysis (BIA). A BIA helps identify the most critical business operations and the dependent resources such as people, technologies and assets required to maintain routine operations. As part of the BIA, desired recovery time and recovery point objectives will be identified to help understand the impact a major disruption to the business process would have on productivity, revenue, and customer satisfaction [1].

The planning and implementation of a successful disaster recovery strategy can be quite complex and highly customized to a firms' business natures and other unique situations. For most small and medium size businesses, budget constraints often preclude the adoption of a sophisticated disaster recovery plan that offers promising recovery time and fool-proof recovery processes. In reality, most challenges related to DR planning are related to the complexity of communication paths and failure points [2]. Clearly, a complex communication topology will result in a higher failure rate. While complexity is somewhat related to system failure, any DR plan should be assessed by looking into the human resources cost associated with personnel support. In today's world, hardware is relatively cheap and creating complex models with very high degrees of fault tolerance is possible with even limited funding. However, devising, configuring, maintaining and understanding data management requires a substantial personnel commitment. Because personnel is typically the greatest cost within the IT budget, understanding model complexity can be a very useful tool in assessing the potential personnel costs related to any disaster recovery model. The key is determining a way to assess individual recover sites and network complexity in relation to the number of hosts deployed [3].

Using the *scale-free degree distribution* formula, the authors present a methodical discussion about the network characteristics of various disaster recovery options [4]. Based on some disaster recovery models for computing domains, the authors discuss the investment required for technology infrastructure and personnel support. Further, instead of focusing on expensive options that are beyond small and medium sized businesses' DR budgets [5], this paper proposes the adoption of a modified hot site that leverages commercial hardware using standard PC components. To the best of our knowledge, this paper is the first work that employs the above formula in assessing complexity of communication paths and failure points for recovery sites. Some important implications of this paper include devising a practical assessment tool for DR planning, recovery investment analysis through comparisons of recovery sites spending, and devising infrastructure requirements for various recovery options.

AVAILABLE OPTIONS

No organization is totally immune to catastrophes, and the danger of centralizing all data and logistics in a single location is often not heeded and results in huge risks which can result in very costly consequences if the company cannot continue to function after a disaster. To tackle this problem, it is advisable to devise plans that utilize available technologies to regain access to vital information systems (IS) in a reasonable time at an affordable cost. DR plans can vary greatly in sophistication and investment cost. The simplest plan might simply involve a tape backup, while the most complex might feature fault tolerance on multiple levels and might be similar to the plan offered by Wang et al. [6]. Organizations have to devise a cost effective, easily deployed DR plan to ensure business resiliency. Such a plan must meet feasibility, consistency, reliability, and specific IT environment characteristics. These are the important issues related to whether or not a DR plan is feasible in terms of human

resources, technology infrastructure and recovery time [7].

Recovery Sites Characteristics

Currently, there are three commonly available data backup site categories: cold, warm or hot sites. Figure 1 is a graphical representation of the various available options in regard to recovery expenses (investment) versus recovery time. Infrastructure complexity and investment cost are the two major considerations in choosing an appropriate DR model. For small and medium-sized businesses, the availability of technologies and size of the DR budget are limitations that dictate reasonable recovery spending. However, even with resource limitations well thought out models can still yield effective benefits. Therefore, decisions related to allocating these resources are important in ensuring organizational sustainability. Bryson et al. [8] advocates using mathematical modeling in analyzing and designing DR models. Past research indicated that the more physical components (hosts) in the DR infrastructure, the greater the probability of a hardware failure. Certainly, additional hardware can provide additional fault tolerance, but it will also increase the DR expenses particularly from a personnel perspective.

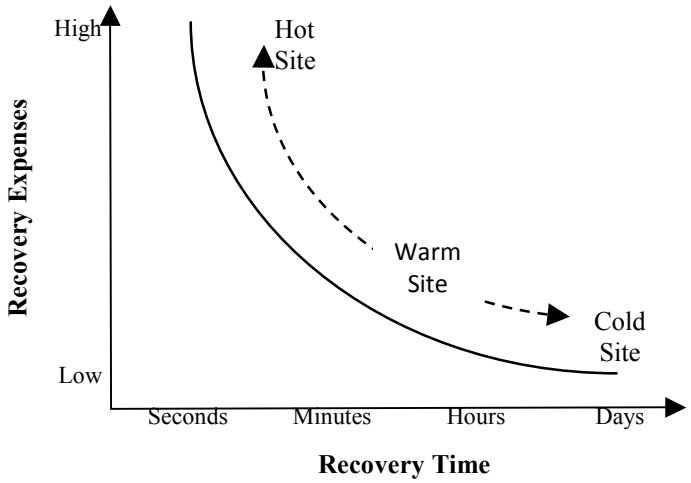

Figure 1: Recovery Expenses Versus Recovery Time

Building on the existing research, we employ the scale-free degree distribution theory to determine network growth by applying the formula N*(N-1)/2 to assess the complexity of the communication path within each DR model [9]. Our assumption rests on the premise that a complex model will result in more possible failure points and be more difficult to support from a personnel perspective. The formula allows us to discover the complexity of communication paths and possible failure points based on the total number of computers in production and the configuration of the replication process. In a simple client/server model that typically consists of a client computer, a server computer, and a network (LAN) connecting them together (see Table 1) the complexity index would be one and result in three total model failure points. In this example, a network complexity of one reflects a very simple model. The number of individual possible failure points is three, which is arrived at by adding the number of computers (N) to the network complexity (C). Because this backup scenario is applied to only one host (one instance), the model's total number of failure points is still three.

For the rest of the computations presented in this paper, we based our assumptions on a six-host production model because such a quantity is regarded as a representative model for most small and medium sized businesses. Furthermore, that is the number of production hosts in our computing domain and we had experience working with that number. However, any number of hosts (computers) might occur. It is common for organizations to have separate hosts for various applications such as accounting and inventory to manage security and performance indicators. Further, additional hosts are often required to support networking activities such as world wide web (WWW), domain name service (DNS) and remote file systems.

Table 1: Disaster Recovery Models Complexity and Failure Point Computation

Description	Basic	Cold	Warm Site		Hot Side		Modified Hot Site	
	Local	Local	Production	Backup	Production	Backup	Production	Backup
Assumptions	N(L)=1 I(L)=1	N(L)=1 I(L)=6	N(P)=6 I(P)=1, W=1	N(B)=2 I(B)=6	N(P)=18 I(P)=1, W=2	N(B)=12 I(B)=1	N(P)=3 I(P)=1,W=2	N(B)=2 I(B)=1
Complexity C=N*(N-1)/2	2*(2-1)/2 = 1	2*(2-1)/2 = 1	6*(6-1)/2 = 15	2*(2-1)/2 = 1	18*(18-1)/2 =153	12*(12-1)/2 =66	3*(3-1)/2 =3	2*(2-1)/2 =1
Failure Point/ Instance FP= N+C	2+1 =3	2+1 = 3	6+15 = 21	2+1 = 3	18+153=171	12+66=78	3+3=6	2+1=3
Total Failure Point I*FP	1*3=3	6 Instances 6*3=18	1*21=21	6 Instances 6*3=18	1*171=171	1*78=78	1*6=6	1*3=3
Model Failure Point FP(P)+ P(B)+ FP(W)	NA	NA	21 + 18 + 1 =40		171 + 78 + 2 =251		6 + 3 + 1 =10	

Note: *N=Computer; I=Instance; L= Local Site; P=Production Site; B=Backup Site; W=WAN; NA: Not Applicable*

Cold Site

A cold site is typically the most inexpensive back-up option to operate and it involves minimal cost to set up. There are no functioning backup copies of the data at the primary data center, and often no additional hardware is required if the tape backup systems are already available. The recovery methodology used is essentially a restoration from tape, to hardware at a remote site with a daily update. A basic visual description of a cold site is provided in Figure 2. Due to its simple form in the DR setting, cold site back-up options require additional time to recover data following any disaster. Each host is backed up independently. Hence, for any one backup option, only a host and a tape backup storage device are required. Most small and medium sized businesses following our example would be required to backup six hosts (i.e., computers).

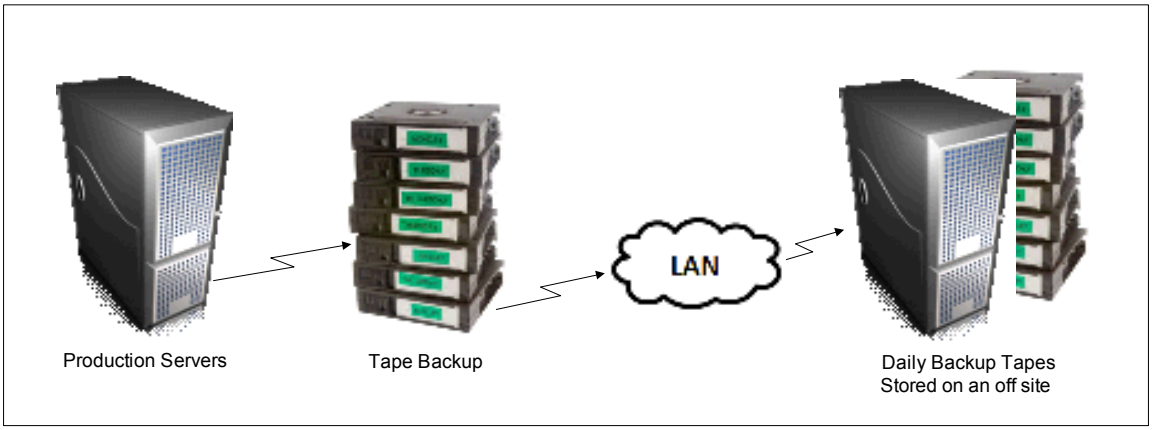

Note: There are six instances of this basic drawing because there are six independent hosts .

Figure 2: Cold Site Disaster Recovery Model

As shown in Table 1, the scale-free degree distribution formula can be applied to discover the complexity of a cold site model of a host/tape backup pair, 2*(2-1 host)/2 = 1. Therefore, the network complexity for this one instance equals one. When adding this value to the number of hosts required, the possible individual failure points equal three, similar to the standard simple client/server example delineated earlier. However, in a six-host model, all six computer pairs are backing up independently and thus each host pair's failure points must be accounted for separately. Thus, the resulting number of network failure points is the product of the total number of instances (six hosts) and individual failure points (3). Further, because this is a local model (although the tapes themselves may be sent off site) there is no WAN connection.

Warm Site

A warm site is an alternate location where data could be retrieved after disruption. It is equipped with hardware similar to the primary site but does not store exact copies of the data. Often the updates take place hourly. A warm site is moderately expensive to operate and the cost largely depends on the desired speed of recovery. It may or may not have the same technological capacity as the primary site, depending on the recovery time objectives (RTO) and recovery point objectives (RPO). After a catastrophe, data will have to be restored onto available equipment at this site to resume business operations. A basic graphic depiction of this model is shown in Figure 3.

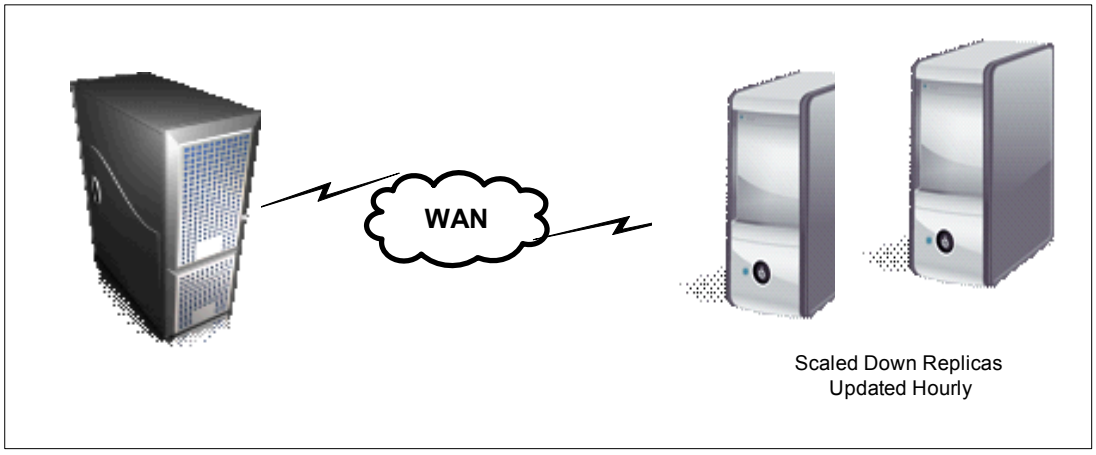

Note: There are six instances of this basic drawing because there are six independent hosts .

Figure 3: Warm Site Disaster Recovery Model

Using the base assumption of a six-host production model, our warm site model would require six local and 12 remote hosts, whereby there are two remote replicas for each production host. The logic is that in the event of a disaster at the production site, one replica in each pair would become the production host which would still leave one replica for backup. The production hosts are configured to share data for backup and performance purposes because in a distributed network, all data required by a given host will not reside entirely on a production host itself. Data are transferred across all hosts and thus it is crucial that all hosts have the ability to communicate with each other. As a result, any given replica in the replica set must offer a high degree of flexibility and fault tolerance. In essence, this is an extension of the local network area (LAN) replication using the RAIH (redundant array of inexpensive hosts) logic but is expanded across a wide area network (WAN) to provide remote replication in case the LAN site and its hosts are compromised.

Since each site is designed to mirror the others, the typical warm-site model would require the production site to have host interaction capability, and the remote site would be equipped with a series of independent connections. This structure allows great improvement in recovery granularity but is very sensitive to the timely data update process due to WAN dependency as well as the back-up synchronization process. The major advantage of this option is higher fault tolerance and better recovery

outcomes, but the recovery investment is also significantly higher than in the cold model.

As shown in Table 1, at the production site the formula would look like: 6*(6-1)/2 =15 + 6 hosts =21 (total unique communication paths among six computers plus the six computers themselves results in 21 possible failure points). The remote site analysis of the replica pairs for complexity yields: 2*(2-1)/2 =1 + 2 hosts = 3 (a two computer replica pair for each of the hosts at the main site, hence 6 instances of these pairs yields a complexity of 18). Therefore, the total complexity is 21(main site) + 18 (remote site) +1(for the WAN connection) = 40. In other words, in the total warm model depicted, there are then 40 possible failure points.

Hot Site

A hot site is the most expensive DR option with full technological capacity that enables a seemingly foolproof recovery processes. Due to its sophisticated information technology (IT) infrastructure, hot sites allow real time synchronization between the primary and alternate back-up site, allowing a complete mirroring of the original data using wide area network links (in our case two independent leased links) and advanced software. Following a disruption to the primary location, the data processing quickly can be relocated to the hot site with minimal loss to routine operations [9]. In any commercially operated hot sites, it is assumed that full connectivity within each site can be achieved within seconds after a disaster. This recovery option in its basic form is depicted in Figure 4.

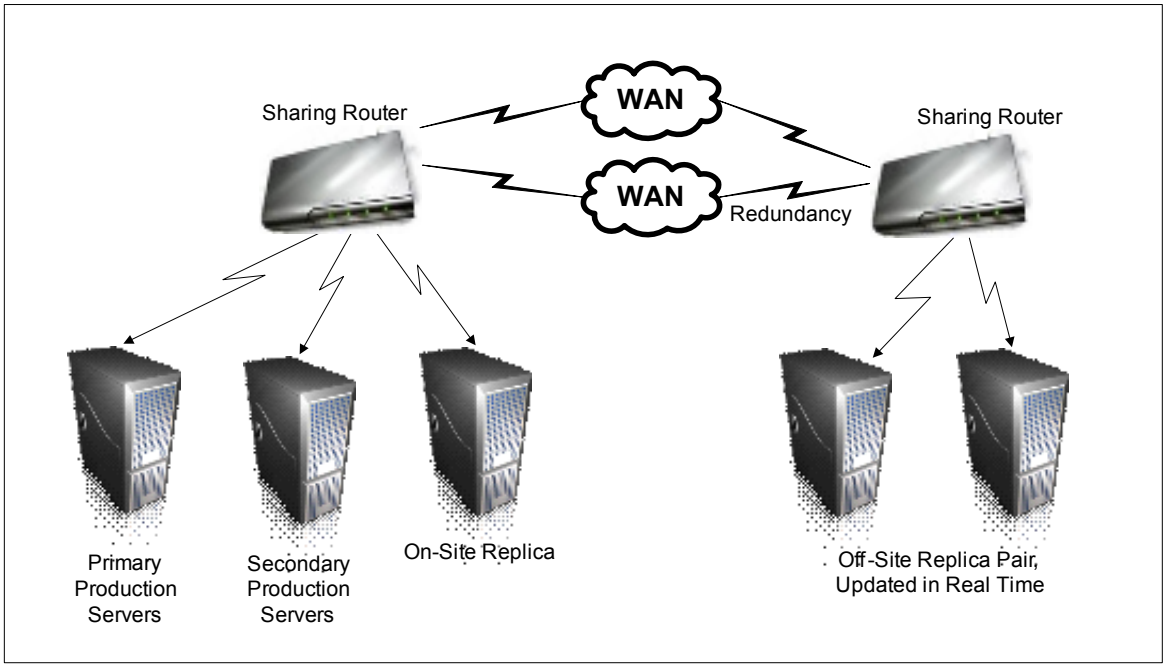

Figure 4: Hot Site Disaster Recovery Model

While a hot site model offers the best fault tolerance and recovery times, it is the most expensive and sophisticated in terms of technology infrastructure. Most likely, it is also beyond the budget allowance of most small and medium sized businesses. In this model fault tolerance has been maximized and recovery granularity minimized. In most hot site options, a RAIH model (with a RAID also employed on the disk level) is used to maximize flexibility and fault tolerance. At the main site each host is replicated twice, it is also replicated twice at the remote site. Hence, if a main host is compromised, there are twice as many chances for it to recover locally and remotely. This means if the entire main site is lost, the network has twice as many chances to recover at its remote site. Further, the main and remote site (it is possible to have more than one remote site at added cost, but

our model doesn't for the sake of simplicity) must be connected via a high speed wide area network (WAN) connection(s) so that the remote replicas are updated in real time. Even if the entire main site is lost, little or no data would be lost.

As shown in Table 1, based on a six-host assumption, the local site analysis indicates that each of the six production hosts would require two replicas, hence requiring 18 computers (six hosts * three computers= 18). So the network complexity for the production side would appear as: 18*(18-1)/2 = 153 (path complexity), plus the 18 hosts = 171 (total failure points at the production site). At the remote site, we have to replicate 6 hosts twice each, hence there will be 12 hosts (there is no main host here; it is at the main site). Thus, 12*(12-1)/2 = 66 (path complexity), plus the 12 hosts would result in 78 for the remote site. The total model complexity is 171(main site) +78(remote site) +2(WAN connections) = 251 possible failure points.

Modified Hot Site

A modified hot site is a recovery option that provides partial benefits of a hot site with a lower DR investment. We propose this option based on our success in leveraging the benefits of host virtualization via creating multiple logical computers (partitioning the resources of one physical computer into six virtualized resource sets) in one single physical computer. Because all production hosts are virtualized into one physical host, this option generates a smaller complexity index and, as a result, has fewer failure points. A modified hot model is an attractive alternative for most small and medium size businesses that can not afford to adopt a hot site option. While the modified-hot side option will not offer the recovery granularity of a true hot site, its intrinsic performance is stable and within acceptable boundaries. Most importantly, this option requires minimum investment and thus it is a very cost effective alternative. Instead of six physical hosts in a traditional warm or hot site, only one physical host is required to house the original six hosts. Thus, all data resides on one single machine. Using a common analogy, this is similar to putting "all your eggs in one basket." However, the logic is to have multiple "baskets" to mitigate risk.

Hence the one now virtualized physical host (containing six logical hosts), is also replicated twice at the main site (of course each replica contains six virtual hosts too). The complexity data for this model follows: 3 (physical hosts) *(3-1)/2 = 3(network path complexity), plus the 3 hosts (physical computers) = 6. The remote site, as before, would require only 2 replications of the main production computer at the main site. This would be calculated as follows 2(physical hosts)*(2-1)/2 = 1(network path complexity), plus the 2 hosts = 3 (total remote site complexity). There is one internet WAN connection Therefore, the total model complexity is 6 + 3 + 1 = 10.

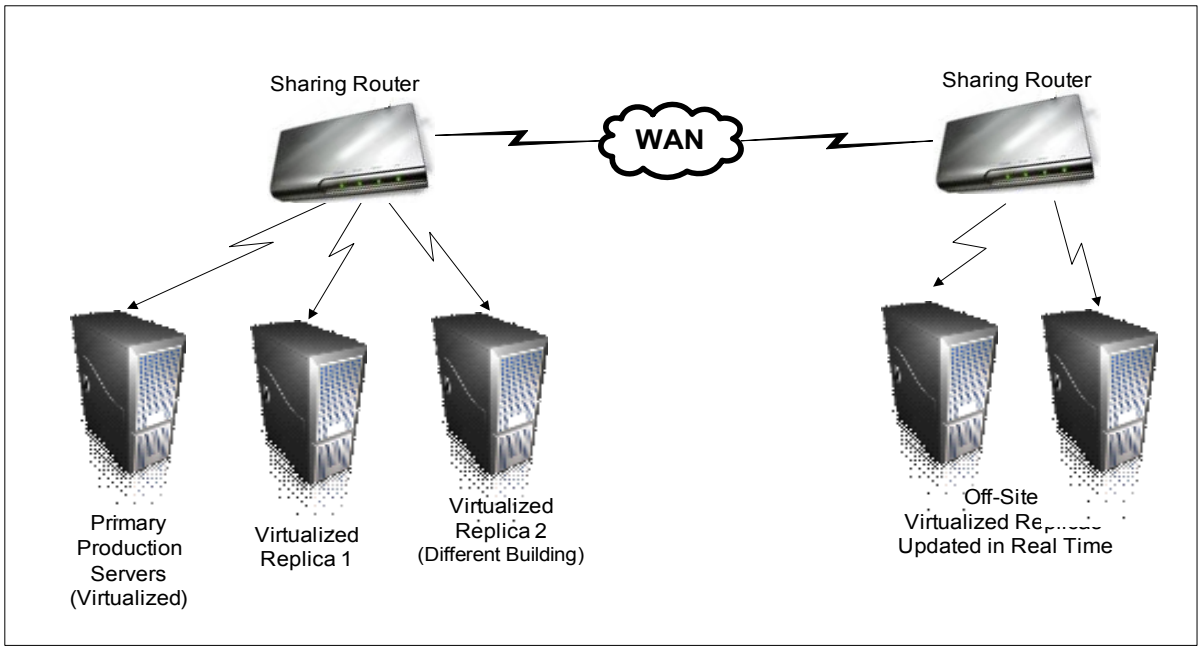

Figure 5: Modified Hot Site (Virtualized) Disaster Recovery Model

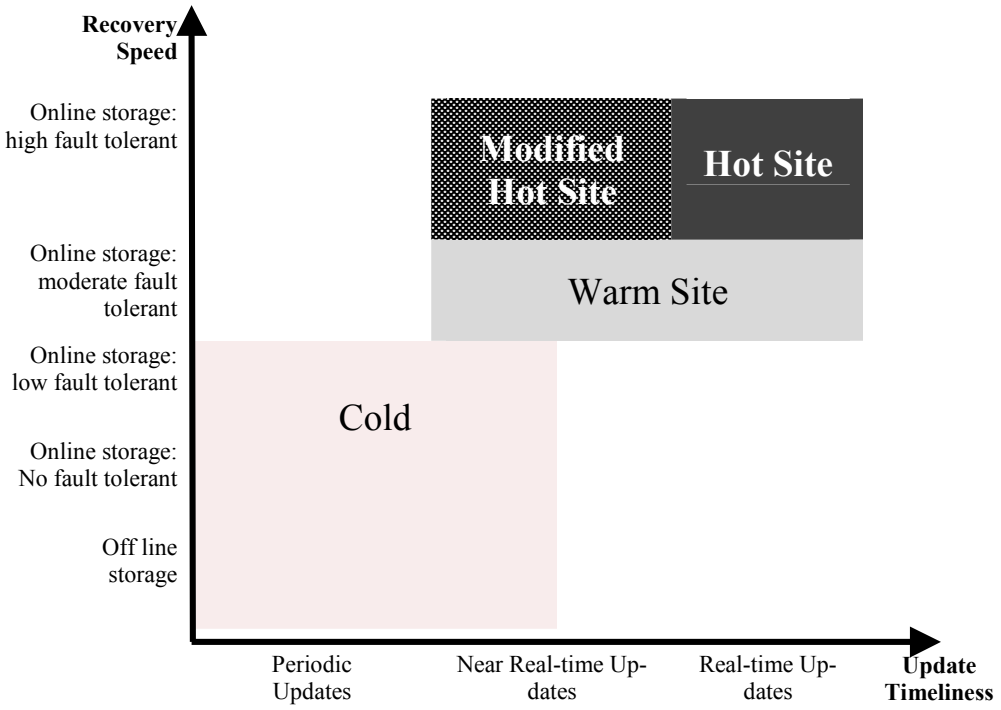

Figure 6: Recovery Speed and Update Timeliness

INFRASTRUCTURE REQUIREMENT

A starting point for selecting recovery site options is to understand the infrastructure required and its associated investment and annual maintenance costs. The basic infrastructure requirement consists of hardware, software and personnel costs. Hardware includes computers, networking equipment, and secondary storage equipment such as tape drives. Each host replication contains daily backups and/or is synchronized to disk as close to real time as cost permits in case of a single geographically localized emergency such as a fire or flood. In case of a multi-point failure such as nuclear-war or an asteroid striking the earth, multipoint offsite backup media would be used to restore data to the new server/s (or switch to already existing replicas in the case of the hot model) and, hence, models that incorporated more than one remote site would better handle these emergencies.

All the above plans require keeping a mirror image of the data for later use in the recovery process as shown in Figures 2-4. Depending on the back-up storage devices, the number of mirror images and effectiveness of the plan, recovery speed can vary remarkably. Sophisticated plans, as presented by Abhang and Chowdry [11], feature multiple image back-ups and quick recovery times. However, these models require expensive and highly advance equipment such as SANs (storage area networks) or other complex storage devices. Generally speaking, higher complexity may result in higher reliability, but will also incur higher technology infrastructure and personnel costs. Instead of focusing on expensive options that are beyond small and medium sized businesses' DR budgets, we propose the use of a modified hot site as a practical solution that relies on commercial hardware and standard PC components (Tables 2 and 3).

Table 2: Disaster Recovery Models For A Six-Host Model

Model	Synchronize Time	Recovery Time	Back-up Site Characteristics	No. of Computers	Tolerance Support
Cold	Days	>24 Hours	Off site backups	12	Limited
Warm	Hours	1-24 hours	Limited physical mirroring	18	Moderate
Modified Hot	Minutes	1 hour	Virtual mirror image	6	High
Hot	Seconds	Minutes	Physical mirror image	30	Very high

Note: There is no remote backup site in the cold model. Hence, the total no. of hosts = 2 computers per instance.

Table 3: Disaster Recovery Investment and Annual Maintenance Costs Based On Six Instances

Recovery Sites	Software: Server Site	Hardware: Remote Site	Bandwidth Cost (WAN)	Personnel Cost	Total Cost:
Cold	2,500	6,000	0	10,400	18,900
Warm	5,000	6,000	4,800	26,000	41,800
Modified Hot	0	3,000	1,200	13,000	17,200
Hot	50,000	120,000	120,000	52,000	342,000

Network Infrastructure

The network infrastructure of a DR site involves line speed, topology, how the connections are laid out, and whether the connections are contained in a WAN or a LAN. To limit the scope of this paper, our discussion focuses on these three concepts. In terms of network infrastructure, topology and speed matter. It is important to note that any change made to the production host must propagate to all other replicas of that host or file system. For remote replicas, the update process also requires a WAN connection. The algorithm used in this replication process can have an impact on the complexity of the network infrastructure. Complexity may be beneficial if it increases the degree of fault tolerance, but if adequate WAN bandwidth is not provided the desired granularity will not be obtained. Conversely, a high performance bandwidth investment may not be worthwhile or cost effective when poor model design causes the updating process to be inherently slow causing the desired recovery granularity to become unattainable [6]. Therefore, instead of replication in a WAN environment, small and medium sized businesses should consider replication in a local area network (LAN) as the first line of defense. LANs provide speedy and inexpensive network configurations and can support reasonably sophisticated replication methods. Ultimately, data will need to be replicated remotely. To improve efficiency, remote replica models must incorporate a degree of data stream optimization and better tuning than LAN models. Therefore, compression strategies and only updating the data that actually changes instead of a complete replica copy are paramount to getting the most out of a speed limited WAN link.

Overview of Replica Strategies

To gain some understanding of the complexity, it may be useful to examine a common fault tolerance concept, the RAID (redundant array of inexpensive disks) logic. In the proposed modified hot site option the array concept can be expanded to the host level, and the RAIH (redundant array of inexpensive hosts (computers)) concept might be more appropriate because we are mirroring hosts instead of just disks. In a six-host model, all six computers of the production site are equipped with computing power for an instructional domain. These six hosts perform the following functions: host DNS (domain name service), maintain a global file system, enable website service, allow email communication, serve as a firewall, and provide instructional support. The capability of RAIH allows the data to support all functions to be replicated across all hosts, and therefore the data is available for any given function when needed to support its primary purpose. In a six-host model, each local disk will have to be

logically divided into six separate partitions. Although this model offers excellent fault tolerance, it is inherently complex, and sophisticated personnel are required for system support. In a LAN environment, the additional complexity and inter-processor communication might be practical because there is adequate bandwidth to support them. However, in a WAN environment, the additional communication overhead may negate the model due to the network's inability to complete the needed updates timely and cost effectively due to the slower speeds and the additional number of bytes that a RAIH would generate.

Software Costs

In regard to software, the costs vary based on the complexity of the model and whether shareware or commercial software is used. For the cold site, a commercial tape back up package with 6 licenses was selected. In the case of the warm model, a commercial backup/replica package with 6 licenses was chosen. Since the modified warm model was devised from scratch using the Unix operating system the openness allowed the use of shareware software at no cost. The complexity and high reliability needs of the hot model necessitated an enterprise level backup replica package with unlimited server licenses.

Hardware Costs

As would be expected, the variance in complexity led to variation in the hardware needs. In all cases we are assuming the production side hardware was in place, as well as the networking infrastructure minus the WAN bandwidth costs. For the cold model, six low end PCs with tape drives at $1,000 each (this allows for simultaneous backup of all six hosts) were required. In the warm model, six mid range server level hosts (at $1,000 each) were needed. For the modified hot virtualization model the hardware needs were reduced to three mid range server level hosts at $1,000 each. Lastly, as expected, the hot model required the most hardware, 12 high range enterprise server level hosts at $10,000 each.

Bandwidth Costs

In regard to band width, the tape back up equipment would reside on site so there are no WAN bandwidth requirements for the cold site. A traditional warm site would rely on leased lines for the sake of security. In this case a leased line (point to point) at 300 miles at 12mbs was selected to meet the bandwidth requirements. The modified hot site developed to save money risks using the internet to provide connectivity (a VPN is used to enhance security) and an internet cable connection at 10mbs has been selected based on the data observed in [12]. To provide the massive bandwidth and reliability required by the hot site dual 40 Mbs leased lines at 300 miles were the configuration chosen.

Personnel Cost

Lastly, the personnel costs also varied greatly from model to model. In the cold model the primary operational costs were tape backup operations and the personnel cost is estimated at 10 hours a week at the rate of $20 an hour. In the warm model, we estimated that 10 hours a week, at $50 an hour for system/network engineering personnel, would be required. The lower complexity of the modified hot model reduces the personnel needs down to 5 hours a week (due to reduced hardware) at $50 an hour for system/network engineering personnel. The hot site had the greatest personnel needs estimated at 20 hours a week (due to added hardware) at $50 an hour for system/network engineering personnel [13].

DISCUSSION

Disaster tolerance, as previously stated, is the ability to maintain ongoing productive operations even in the face of a catastrophe. This is an important consideration since high availability is achieved by providing redundant components; if one fails, another part is still available to do the job. To better understand the key issues in selecting a recovery site, it is helpful to identify the complexity of a chosen model's communication paths and its possible failure points. In reality, the most challenging goals surrounding DR planning are related to deciding on the appropriate number of hosts, degree of fault tolerance desired, appropriate granularity and attainment of all of those goals within the available budget. Applying the scale-free degree distribution formula, this paper demonstrates how a complexity index of various recovery sites can be computed and how to identify network failure points. The paper also proposes the adoption of a modified hot site for its attractive cost effectiveness and many inherent benefits similar to those of a hot site option.

A key advantage for the proposed modified hot site is cost effectiveness. By deploying virtualization to reduce the number of physical hosts and using shareware software, firms can develop a structured and actionable DR plan that attains many of the benefits of a hot site model. The application of virtualization is a formal approach to DR planning which enables effective DR solutions that are less complex, more cost effective, and close to the performance level of the hot model. We assert that our proposed option is an optimal solution for small and medium sized businesses due to its capabilities to enable

firms to simplistically map out the dependencies between critical business processes, people, IT assets, and other resources. It can also perform simultaneous functions in hosting DNS, maintaining a global file system, enabling website service, allowing email communication, serving as a firewall, and providing instructional support all in a single physical host, while still maintaining the separation of those services for performance and security purposes [14].

The models described herein are certainly not the only models available. In fact, new models could easily be devised taking pieces from the basic models described herein. Each company will vary in regard to the characteristics of the original production configuration and their data recovery goals. While the models presented herein may not be directly transferrable to another company's situation the process of evaluating potential models is transferrable. Core to our analysis in this paper was the scale-free degree distribution algorithm. It can be very useful in assessing model complexity thereby providing some idea of the fault tolerance provided in relationship to the personnel and hardware required. In our case we suspected that virtualization would be beneficial, but not to the degree observed in our evaluations herein. Specifically, the evaluation process we used to assess model complexity, cost and granularity confirmed it was adequately robust and a very cost effective choice.

REFERENCES

[1] Balaouras, S. "Market Overview: Business Continuity Planning Software," *Forrester Report*, May 30, 2007.

[2] Hill, J. "Business Continuity: Implementing Disaster Recovery Strategies and Technologies", *Aberdeen Benchmark Report*, March, 2008, p.1-19.

[3] Moe, T. L, Gehbauer, F., and Senitz, S. "Balanced Scorecard For Natural Disaster Management Projects," *Disaster Prevention and Management*, Volume 16, Number 5, pp.785-806.

[4] Albert, R., and Barabási, A-L. "Dynamics of Complex Systems: Scaling Laws for the Period of Boolean Networks" *Physical Review Letters*, Volume 84, Number 24, 2000, pp. 5660.

[5] Balaouras, S. "Building the Business Case For Disaster Recovery Spending," *Forrester Report*, April 3, 2008, pp.1-16.

[6] Wang K., Su, R-D., Li, Z-X., Zhen, C., and Zhua, L-H. "Robust Disaster Recovery System Model," *Wuhan University Journal of Natural Sciences*, Volume 11, Number 1, 2006, pp.170-174.

[7] Balaouras, S. and Schreck, G. "Maximizing Data Center Investments For Disaster Recovery and Business Resiliency," *Forrester Report*, October 5, 2007, pp.1-13.

[8] Bryson, K-M., Harvey, M., Joseph, A., and Mobolurin, A. "Using Formal MS/OR Modeling to Support Disaster Recovery Planning," *European Journal of Operational Research*, Volume 141, 2002, pp. 679-688.

[9] Baccaletti, S., Latora, V., Morento, Y., Chavez, M. and Hwang, D.U. "Complex Networks: Structure and Dynamics," *Physics Reports*, Volume 424, 2006, pp.175-308.

[10] Yamato, J., Kan, M., Kikachi, Y., Takaya, M., Tomi, M. and Adachi, T. "Outline of Disaster Recover Architectures," *NEC Technical Journal*, Volume 1, Number 4, 2006, pp. 29-32.

[11] Abhang, S. and G. Chowdry. "WDM-Based Storage Area Networks for Disaster Recovery Operations," *International Journal of Computer, Information, and System Science and Engineering*, Volume 1, Number 4, 2007, pp. 493-495.

[12] Guster, D. C., Safonov, P. I., Hall, C., and Sundheim, R. "Using Simulation to Predict Performance Characteristics of Mirrored WWW Hosts," *Issues in Information Systems,* Volume 4, Number 2, 2003, pp.479-485.

[13] Krojnewski, R., and Nager, B. "Disaster Recovery: It's Not Just An IT Problem," *Forrester Report*, November 13, 2006.

[14] Guster, D. C., McCann, B. P., Kizenski, K. and Lee, O.F. "Cost Effective, Safe and Simple Method to Provide Disaster Recovery for Small and Medium Sized Businesses," *Review* of *Business Research*, Volume 8, Number 4, 2008, pp. 63-71.

Note: This paper is based on and is an enhancement of a paper submitted for presentation at the 2009 AMCIS conference.

AUTHOR BIOGRAPHIES

Dr. Dennis Guster is a Professor of Computer Information Systems and Director of the Business Computing Research Laboratory at St. Cloud State University, MN, USA. His interests include network design, network performance analysis and computer network security. Dennis has 25+ years of teaching experience in higher education and has served as a consultant and provided industry training to organizations such as Compaq, NASA, DISA, USAF, Motorola, and ATT. He has published numerous works in computer networking/security and has undertaken various sponsored research projects.

Dr. Olivia F. Lee is an Assistant Professor Marketing at St. Cloud State University, MN, USA. She has worked as an operation manager at two university hospitals, and as a senior e-business market analyst in a business-to-business company prior to her academic career. Her research work focuses on technology practices in business environment, service organization and business resilience strategy. She has published her work at Psychology and Marketing, Healthcare Marketing Quarterly, International Journal of Organization Analysis and Review of Business Research.

Brandon McCann is a graduate student at St. Cloud State University, with over 6+ years of experience in the IT environment. Brandon has worked as an Information Security Consultant and is a Microsoft Certified Professional.

Dr. Mark B. Schmidt is an Associate Professor of Information Systems and the Interim Director of the Center for Information Assurance Studies at St. Cloud State University. He has works published in the several Communications of the ACM, Journal of Computer Information Systems, Journal of End User Computing, Journal of Global Information Management, Journal of Internet Commerce, Mountain Plains Journal of Business and Economics, and the International Journal of Information Security and Privacy. His research focuses on information security, end-user computing, and innovative information technologies.