Components of Disaster Tolerant Computing

Chad M. LAWLER
Data Return, LLC
Irving, TX 75039, U.S.A.
netdeveloping@hotmail.com
and
Michael A. HARPER
Critical Infrastructure Protection Center, SPAWAR Systems Center Charleston
Department of the Navy
North Charleston, South Carolina 29419, U.S.A.
harper.michael@gmail.com
and
Mitchell A. THORNTON
Department of Computer Science and Engineering, Southern Methodist University
Dallas, Texas 75275, U.S.A.
mitch@engr.smu.edu
SMU School of Engineering,
Box 750335, Dallas, TX 75275-0335

ABSTRACT

This paper provides a review of the components of disaster tolerant computing and communications and reviews the current state in light of recent man-made terrorist events. The paper examines the relationships between disaster tolerant systems, Information Technology (IT) application availability and executive level management visibility necessary for successful system operations in the event of a catastrophic disaster; one which causes rapid, almost simultaneous, multiple points of failure in a system, as well as a single points of failure that escalate into wide catastrophic system failures. The technology, process and human resource challenges of traditional disaster recovery approaches to disaster preparedness are outlined. The risks of IT application downtime attributable to the increasing dependence on critical information technology applications operating in distributed and unbounded networks are explored. A general method for disaster tolerance is proposed which mitigates unplanned downtime through a disciplined approach of IT infrastructure design based on redundancy and distributed components with special attention given to the ability of executive level management to comprehend the value of uptime of an application and make appropriate capital investment. The importance of executive visibility into the system wide impact of downtime and the resultant effects on the costs of downtime of critical systems is explored.

***Keywords**: Disaster Tolerance, Disaster Tolerant Computing and Communications, Survivability, Application Downtime, Disaster Recovery, DR, Business Continuity Planning, BCP, Executive Visibility, Cost of Downtime, Value of Uptime*

1. INTRODUCTION & BACKGROUND

A disaster is an event that can cause system-wide malfunctions as a result of one or more failures within a system which may be caused by a single-point failure or by a plurality of single-points of failure that occur either simultaneously or nearly simultaneously in a temporal sense by either a man-made or natural event. A catastrophe can occur as the result of the occurrence of a disaster and may be avoided by using disaster avoidance mechanisms [1].

Fault tolerant system design is a mature discipline based on reducing risks and impacts associated with single points of failure in a system. Disaster Tolerance (DT) in computing and communications systems refers to the ability of IT systems, communications infrastructure, and business or organizational processes that depend on these systems, to maintain functionality throughout the occurrence of a disaster when it has occurred. Disaster Tolerance provides an ability to continue operations uninterrupted despite occurrence of a disaster that significantly interrupts normal organizational operations. Specifically, within DT, critical business functions and technologies continue operations, as opposed to resuming them. Disaster tolerance is a superset of fault tolerance methods in that a disaster may occur which causes rapid, almost simultaneous, multiple points of failure in a system, as well as a single points of failure, that escalate into a wide catastrophic system failures [1].

Businesses and organizations that adopt redundancy-based approaches traditionally rely on Disaster Recovery (DR) techniques to protect critical systems. Disaster Recovery, a subset of Business Continuity Planning (BCP), is a closely related term used to describe methodologies to create and execute a plan for how an organization will resume partially or completely interrupted information technology, organizational, or business critical functions within a predetermined time after a disaster or disruption has occurred. Effective Disaster Recovery and Business Continuity Planning should identify the impacts of the loss of a critical business facility, resource or process in the event of an unplanned with the specific intent of identifying required recovery timeframes and resources. In addition to Service Level Agreements for organizations, recovery management service levels such as Recovery Point Objective (RPO), defined as the amount of data loss that is acceptable, if any, and Recovery Time Objective (RTO), the amount of downtime that is acceptable, if any, are becoming more critical in defining specific recovery metrics [2]. Disaster Recovery and Business Continuity Planning efforts are specifically targeted at reducing operational risk and therefore overlap with traditional risk management practices. DR and BCP commonly utilize IT services and applications along with fault tolerant systems and methodologies to help achieve recovery or continuity [3].

In traditional DR & BCP plans, attention is commonly given to contingencies for natural disasters such as hurricanes, tornados, floods, and earthquakes. However, a disaster may be any event that prevents a business from accessing necessary data and systems to conduct normal business operations. In the past, it may have been acceptable to assign a very low probability to the risk of major disaster occurrence. However, with the rising potential for terrorist activity, this assumption is no longer the case [3].

Models for disaster tolerance differ from those for fault tolerance since they assume that failures can occur due to massive numbers of individual faults as well as a single point of failure. Specifically, the system model can be described as multiple individual system faults that occur nearly simultaneously or close together in time as a series of related events. A naïve way to provide disaster tolerance in a system is to utilize redundancy with replicated components located in geographically disparate locations [1].

However, the approach of systems and infrastructure replication with geographic disparity has six significant general consequences:

1. All elements of a system must be replicated in order for system functionality, including data, servers, storage, applications, Wide Area Network communications, and in particular, human IT resources, which are difficult to replicate.
2. Data replication and synchronization between redundant systems becomes problematic over geographically disparate networks.
3. The complexity of a system increases as the level of redundancy increases, making the components of redundant systems more difficult to manage and complex to maintain.
4. The costs of larger redundant systems are commonly high and discourage capital investment and implementation.
5. Replicated or redundant IT systems commonly implement Disaster Recovery practices to fail over or recover system functionality at a replicated site. With limited success rates.
6. Some systems are so large that it is impractical to replicate them (for example, the United States electric power grid).

2. CURRENT STATE OF DISASTER RECOVERY

The terrorist events of September 11, 2001 and the US Northeast power outage of August, 2003, combined with Hurricane Katrina of 2005, provide recent examples of devastating man-made disasters and massively destructive natural disasters in the US. Some firms affected by the September 11, 2001 attacks, who did have well-developed and thoroughly tested Business Continuity Plans in place were able to recover partial business operations within several days of the terrorist attack. However, many of these businesses have still not fully recovered six years after the event [4].

September 11th particularly serves as a significant event that has dramatically altered the manner in which political organizations approach disaster management. Surprisingly, many areas of the commercial business sector have not demonstrated a similar response. Survey data suggests that top business executives are not focused on Disaster Recovery or Business Continuity Planning. According to survey data from Harris Interactive and SunGard Availability, the majority of US business executives believe their companies are in fact less prepared to deal with a disaster than in years prior to 2005 [4].

Surprisingly, data from studies on September 11 indicates that 9-11 has in fact had relatively little effect on the spending patterns of US mid-sized business on proactive preparedness activities such as protection and security. The Conference Board released a 2005 report on corporate security practices, sponsored by the US Department of Homeland Security, based on a survey of chief executives and other top officers in a wide range of mid-sized US companies (with annual revenues of between $20 million and $1 billion). Data from this report indicates that despite mounting evidence to the contrary, the majority of US mid-market companies believe their business's current spending on security is adequate as a sound business investment that will proactively reduce the risk and impact of a disaster. Many of these companies view these business costs as an expense that should be minimized. [5]

The small percentage of organizations that have the resources available, foresight and capability to consider the risks and costs of mitigating against IT application and business process outages commonly invest in Disaster Recovery and Business Continuity plans. DR solutions traditionally implement alternate 'hot', 'warm' or 'cold' failover sites with varying degrees of IT infrastructure readiness and availability. Unfortunately, such efforts are often made after an IT solution has been designed and implemented, not before, where it could have the most beneficial effect on architecture and appropriate implementation and maintenance [3].

In many cases where replication technology is implemented by business IT organizations, replication and failover process failure is common. Statistics show replication and failover failure has five primary causes.
Secondary failover environments are often not ready for the failover process itself to occur. Manual human error occurs within the failover process. The failover process is dependent on critical knowledge experts who are unavailable during crises. Failover processes are unable to scale in disaster situations where rapid, almost simultaneous, multiple points of failure escalate into system wide catastrophic failures. Finally, and perhaps most common, assumptions made regarding failover are incorrect and result in a lack of successful failover [6].

Such efforts are often unsuccessful in reaching the goal of providing business process or IT application continuance in the event of a disaster. Instead, these efforts attempt to force an application or technology solution to function in a manner in which it was not designed and do not have functionally adequate processes, technology or support resources to enable successful Disaster Recovery failover [6]. As a result, a large portion of capital and resource investment in Disaster Recovery is literally wasted in the failed recovery processes itself, reducing the value of this investment, as it does not produce the desired result: IT infrastructure, applications and business process functionalities that are disaster tolerant [3].

As a result, traditional Disaster Recovery and Business Continuity Planning and practices are often not sufficient to protect businesses and organizations from IT systems and network outages, nor do they enable IT applications and business processes to adequately continue operations throughout the occurrence of a disaster. Traditional DR technologies and practices fail to adequately provide recovery capabilities for organizations to survive major disasters such as the loss of an entire building, city block or large portions of a city itself. In actuality, these practices leave organizations and businesses vulnerable to complete organizational failure in the event of a disaster. Strategy, priority, management, investment, personnel and technology challenges surrounding DR and BCP render these practices ineffective. This existing lack of effective DR and BCP methodologies has crippling potential for businesses and government agencies alike[3].

3. THE COSTS OF DOWNTIME

Executive visibility, with regard to information systems, is defined as the ability of executive management to understand the business aspects of an information system or application and to obtain a comprehensive insight into the financial aspects and contributions of Information Technology systems and applications. Executive visibility should include awareness of service level agreements (SLA) including: corporate compliance and governance, application stability and availability of IT applications and related infrastructure on which the business depends. In addition, executive visibility includes the requirement that key decision makers understand the financial costs of downtime as well as the value of uptime of the application [3].

However, in observing IT operational outages as well as evaluating potential large-scale system or organizational failures, executive managers commonly do not have adequate information regarding the actual financial costs of such downtime and outages. Consequently, the value of uptime is also often not understood. Instead, a lack of visibility into the practical impact of such outages on business processes, customer service, product/service delivery and revenue generation tends to be more common [3].

Eagle Rock Alliance conducted study as a joint effort between Contingency Planning Research, and Contingency Planning & Management Magazine titled "2001 Cost of Downtime". A subset of their findings are as follows: 46% said each hour of downtime would cost their companies up to $50K; 28% said each hour would cost between $51K and $250K; 18% said each hour would cost between $251K and $1M; 8% said it would cost their companies more than $1M per hour [7].

The Northeast power outages of August 2003 cost New York City businesses more than $1 billion, $36 million per hour [8]. Mirifex Systems LLC and the Center for Regional Economic Issues at the Weatherhead School of Management in Ohio lost more than $50,000 per hour of downtime during the blackouts. That adds up to about $400,000 for an eight-hour day [9]. These per-hour loss numbers are in line with a 2000 study conducted by Contingency Planning & Management (CPM) who concluded that companies with annual revenues less than $100 million dollars were likely to experience a $50,000 per hour economic loss from a full interruption in a given year on average. The financial cost per hour loss increased substantially with the annual revenue size of the company. Companies with annual revenues greater than $3 billion were likely to experience an hourly loss rate of more than $1 million per hour. In addition, the study identified those industries on average, experience different per-hour loss rates. Companies in the financial and banking industry experience on average greater hourly loss rates than most other service industries, such as transportation [9].

All too often, the financial cost of implementing redundant applications or hot/warm failover sites prevents management from implementing these technologies. A system that could model and integrate information detailing SLA compliance, cost of downtime, value of uptime, as well as stability and availability statistics could assist in providing greater executive visibility to management staff. This information, in turn would allow management greater insight in making decisions regarding IT applications, infrastructure and business continuance planning. Executive management equipped with accurate information regarding the financial ramifications of application downtime would be able to more readily engage in the cost benefit analysis of implementing an IT infrastructure that is disaster tolerant. With appropriate executive visibility, management would have supporting information to budget for the costs of implementing technology and applications that are able to survive traumatic disruptions. It is essential for executive level management to fully understand the value of uptime of a particular application or IT infrastructure as well as the costs and business impacts of downtime in order to make fully educated decisions regarding disaster tolerant systems and establish business cost benefit justification for capital invest in such systems [3].

4. METHODS

Basic analysis was performed on a simplified, 2-node, geographically distributed architecture that models a secure data transfer network undergoing a failover, emulating the loss of one data site node [3].The system under analysis is structured as a simple series-parallel configuration demonstrated in Illustration 1.

Illustration 1 – Simple 2-Node Failover Model

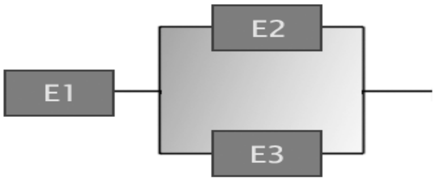

An analytical approach is conducted to determine a Time to Failover model, random variable *t,* for a simplified, series-parallel reliability architecture. A simulation was performed utilizing EMC^2 Legato RepliStor replication and failover software and modeled the time to failover to an alternate data site in response to a disaster and the loss of a data site node. Failover was dependent on manual execution, DNS replicated changes, LDAP infrastructure and replication/failover application communications. Statistical analyses revealed that each element (E_i) of the system follows an exponential distribution of time to failure $T_i \sim \mathcal{E}(\Theta_i)$ for i = 1, 2, … n [3].

Based on data from the simulation results, the following analysis of the system was determined:

The Reliability of the system, R(t)
The instantaneous failure rate, h(t)
The cumulative failure rate, H(t)
The Mean Time to Failover (MTTF)

The following assumptions were made for the system model:

1. Asynchronous data replication between data site nodes
2. DNS replicated changes, LDAP infrastructure were functional
3. Perfect failure sensing and switching
4. Zero failure rate during standby
5. Independent elements
6. Element time to failure is exponential with parameter λ

The system reliability for the configuration is:

$$R_s(t) = e^{-(\lambda_1+\lambda_2)t} + e^{-(\lambda_1+\lambda_3)t} - e^{-(\lambda_1+\lambda_2+\lambda_3)t}$$

Where system mean time to failure is

$$\text{MTTF} = \frac{1}{\lambda_s} = \Theta_s$$

MTTF of the series parallel configuration is determined through the relationship:

$$\text{MTTF} = \int_0^\infty R(t)dt = \frac{1}{\lambda_1+\lambda_2} + \frac{1}{\lambda_1+\lambda_3} - \frac{1}{(\lambda_1+\lambda_2+\lambda_3)}$$

The instantaneous failure rate of the 2-node failover system model, *h(t)*, is calculated through the relationship of the failure density function and reliability the function [3].

$h(t) = \dfrac{f(t)}{R(t)}$ where $f(t) = -\dfrac{d}{dt}R(t)$, which yields:

$$h(t) = \frac{(\lambda_1+\lambda_2)e^{-(\lambda_1+\lambda_2)t} + (\lambda_1+\lambda_3)e^{-(\lambda_1+\lambda_3)t}}{e^{-(\lambda_1+\lambda_2)t} + e^{-(\lambda_1+\lambda_3)t} + e^{-(\lambda_1+\lambda_2+\lambda_3)t}}$$
$$- \frac{(\lambda_1+\lambda_2+\lambda_3)e^{-(\lambda_1+\lambda_2+\lambda_3)t}}{e^{-(\lambda_1+\lambda_2)t} + e^{-(\lambda_1+\lambda_3)t} + e^{-(\lambda_1+\lambda_2+\lambda_3)t}}$$

This allows us to derive the Cumulative Failure Rate

$$\text{H(t)} = \frac{1}{t}\int_0^t h(t)dt$$

$$\text{H(t)} = \frac{1}{t}\int_0^\infty \left[\frac{(\lambda_1+\lambda_2)e^{-(\lambda 1+\lambda 2)t} + (\lambda_1+\lambda_3)e^{-(\lambda_1+\lambda_3)t}}{e^{-(\lambda_1+\lambda_2)t} + e^{-(\lambda_1+\lambda_3)t} + e^{-(\lambda_1+\lambda_2+\lambda_3)t}} \right.$$
$$\left. - \frac{(\lambda_1+\lambda_2+\lambda_3)e^{-(\lambda_1+\lambda_3)t}}{e^{-(\lambda_1+\lambda_2)t} + e^{-(\lambda_1+\lambda_3)t} + e^{-(\lambda_1+\lambda_2+\lambda_3)t}} \right] dt$$

Mean Time to Failover (MTTF) value can be estimated using a (1- α) $\bullet$95% Lower Confidence Interval (Θ_L,∞), based on the condition that simulation discontinued after a fixed amount of total time T_c has elapsed [3].

$$\Theta_L = \frac{2T_c}{\chi^2_{\frac{\alpha}{2},2r}} \quad \text{where}$$

$\chi^2_{p,df}$ is the value of x~ χ^2_{df} such that P(X> χ^2_{df}) = p. Five simulations of time to failover were tested and measured in seconds. The results follow: 22, 27, 33, 47, 73

MTTF is estimated by the point estimate

$$\hat{\Theta} = \hat{f}(\hat{\Theta},\hat{\Theta}^2) \frac{\sum_{i=1}^{5} x_i}{5} = \frac{95}{5} = 19 \text{ seconds.}$$

A 95% lower confidence interval on the mean failover is set, providing a measure of potential variation. Appropriate chi-square test values are $\chi^2_{0.05,10} = 18.31$

Therefore, with a 95% lower confidence limit on Θ the MTTF is 10.38 seconds for the series-parallel system configuration. Using the point estimate, $\hat{f}(\hat{\Theta},\hat{\Theta}^2)$ for MTTF and varying the confidence interval, we determine that $\hat{f}(\Theta_L,\Theta_U^2)$ and $\hat{f}(\Theta_U,\Theta_U^2)$ are worst case scenarios, as these points have the highest variability. In a similar manner, $\hat{f}(\Theta_L,\Theta_L^2)$ and $\hat{f}(\Theta_U,\Theta_L^2)$ have the least variability and are therefore more desirable to work towards achieving a MTTF in this area for increased predictability.

This risk avoidance posture affords mitigation against the costs and consequences of unpredictable IT application downtime provides an organization with the ability to analyze, predict, and rationally accept associated risks, as warranted by an application's availability requirements. [3]:

Illustration 2 – Analysis of Simple 2-Node Failover Model

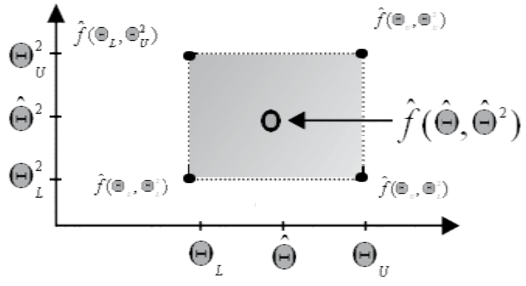
$\hat{f}(\Theta_L,\Theta_U^2)$
$\hat{f}(\Theta_U,\Theta_U^2)$
Θ_U^2
$\hat{\Theta}^2$
Θ_L^2
$\hat{f}(\hat{\Theta},\hat{\Theta}^2)$
$\hat{f}(\Theta_L,\Theta_L^2)$
$\hat{f}(\Theta_U,\Theta_L^2)$
Θ_L
$\hat{\Theta}$
Θ_U

5. DISASTER TOLERANT APPROACHES

The capability to deliver essential services in a constant and continuous manner must be sustained even if a significant portion of a system is incapacitated. In order to achieve Disaster Tolerance, this capability should not be dependent on the survival of a specific information resource, node or communication link. In a wartime environment, essential services might be those required to maintain technical superiority and essential properties may include integrity, confidentiality, and a level of performance sufficient to deliver results in less than one decision cycle of the enemy. Similarly, in the public sector, an IT application system maintaining financial information has the requirement to maintain integrity, confidentiality, and availability of essential information and financial services, even if particular nodes or communication links are incapacitated because of a debilitating event [10].

An effective approach for standard IT solutions development, implementation and support should inherently include an appropriate level of disaster tolerance built into the architecture itself, from initial design through to implementation and management. Unfortunately, due to increased costs and a lack of comprehension of the true costs of IT application or business process downtime, this practice is often not followed. This is particularly true in business scenarios where competition for financial and human resources is intense and is thus often ignored instead of managing the risk of outages and failures. Data indicates that it is in fact common for executive management when faced with capital investment decisions to allocate funds to other areas within a business instead of investing in disaster recovery [110]. Despite recent increases in both man-made and natural disasters, the large majority of businesses and executive management continue not to consider Disaster Recovery or Business Continuity Planning a top priority [12].

Building redundancy and disaster tolerant designs into the initial architecture itself is not a new a concept. However, establishing a proven process that incorporates disaster tolerant technologies early in the IT solution design is different from the current concept of disaster recovery. This approach would alter the way the IT solution design process has historically been done but offers the potential for significant benefit in terms of disaster tolerance [3].

Providing architectural features for disaster tolerance and redundancy into business and organizational IT solutions and applications is a step in a different and potentially financially beneficial direction. However, a disaster tolerant infrastructure and application alone will not solve all of the challenges surrounding the issue of application downtime. In addressing this issue, consideration should be given to technology as well as to the business or organizational strategy and the people and processes that affect the availability of a system. Furthermore, executive visibility into the availability of the application, as well as the resulting insight such visibility would provide, would help technology managers better understand the people, strategy, processes and technologies involved in keeping an application available [3].

Secure and reliable operation of modern IT applications and systems is a standard requirement for both business and government organizations. Due to their large-scale, complexity and non-redundant architecture, modern IT application infrastructures are not commonly capable of tolerating significant disruptions due to man made or natural disasters. As a result, avoiding cascading failures in complex IT application systems is a significant challenge. In standard computer cluster configurations, high availability is often achieved through the use of redundant hardware to eliminate single points of failure. This approach can protect the cluster against hardware faults, such as individual node failures [13].

For some systems, the level of protection a cluster provides is insufficient. An example is systems running financial transaction processing applications where multiple system failures have significant financial impact to the organizations depending on them. For such systems, it is necessary to guard not only against single points of failure but also against multiple points of failure (MPOF). In the case of disasters, this includes protecting against a single massive failure that causes many components to fail, such as the failure of an entire data center which physically contains groups of servers nodes and disk storage subsystems in close geographic proximity [13].

Clusters that are resistant to multiple points of failure or single massive failures require a different type of cluster architecture known as disaster tolerant architecture, which provides the ability to failover to alternate cluster nodes. However, system complexities and people-related processes often render the failover and failback scenarios dysfunctional and inadequate. A superior disaster tolerant architecture involves designing server clusters to share the system load among several geographic cluster nodes in a distributed fashion where the loss of one or more nodes or geographic locations does not significantly impact system functionality [13].

Three specific technologies identified indicate potential for disaster tolerance. A project and technology architecture known as Myriad provides an alternative to data site mirroring for achieving disaster tolerance in large, geographically-distributed storage systems. The Myriad architecture and methodology is implemented through a protocol permitting cross-site checksums which are updated in such a way that data recovery is always possible [14]. An optoelectronic technique which leverages *Dense Wave Division Multiplexing* (DWDM) also holds potential for disaster tolerant technology architecture. This approach employs multiple wavelengths to transmit signals over a single optical fiber allowing system-to-system communication and database replication. Currently, the maximum distance a signal can travel without degradation or decrease in reliability is limited to 100km [15]. Implementation of high availability clustering technologies may also provide information technology infrastructure for disaster tolerance. This technology includes multiple nodes configured in a server cluster that allow simultaneous access to data in a shared file system. Therefore, the view of the file system is effectively the same from any node in the cluster which provides potential for disaster tolerant computing and communication design [15].

A fundamental yet critical step in designing a viable disaster tolerant IT infrastructure and applications is to begin with the idea of disaster tolerance in mind. Disaster recovery and business continuity technologies and plans are often conceived after an application has been designed and implemented, adding into the existing infrastructure disaster recovery functionality features that were not designed into the application itself [3]. Applications and technologies implemented are then intended to function in a manner in which they were not designed. In cases where replication technology is implemented, replication and failover process failure is common [6].

Disaster tolerant applications should be designed from their initial stages with replication, failover, multiple site architecture and other redundant technologies be built into the design itself [3].

6. CONCLUSIONS

The terrorist events of September 11, 2001 and the US Northeast power outage of August, 2003, combined with Hurricane Katrina of 2005, emphasize the need not only to develop disaster tolerant computing and communication systems. The current state of Disaster Recovery and Business Continuity, in light of recent man-made terrorist events, may not be sufficient in their goal of providing business and IT systems recovery. In the event of a disaster, organizational and business continuance is dependent on the ability to continue IT operations and provide continues application and business process availability.

A Disaster Tolerant approach is suggested which has the goal of continuing IT and business process operations, as opposed to resuming them, including redundancy and distributed components with special attention given to the ability of executive level management to comprehend the value of uptime of an application and make appropriate capital investment.

Increased awareness and application of the areas discussed in this paper will provide executive management significant benefit through increased visibility into the business aspects of information systems and applications regarding the value of uptime, the costs of downtime and the associated aspects involved in implementing disaster tolerant IT architectures. A sufficient level of visibility will provide executive level management with the information necessary to make appropriate decisions regarding architecture, implementation, maintenance and support of IT applications for technology infrastructures capable of surviving and adapting to disasters [3].

7. REFERENCES

[1] Szygenda, Stephen A., Thornton, Mitchell A., “Disaster Tolerant Computer and Communication Systems”, Department of Engineering Management, Information and Systems & Department of Computer Science and Engineering, SMU, 2004

[2] R. J. Ellison, D. A. Fisher, R.C. Linger, H. F. Lipson, T. A. Longstaff, N. R. Mead, "Survivability: Protecting Your Critical Systems," *IEEE Internet Computing*, November/December 1999

[3] Harper, Michael A., Lawler, Chad M., Thornton, Mitchell A., “IT Application Downtime, Executive Visibility and Disaster Tolerant Computing”, CITSA, 2nd International Conference on Cybernetics and Information Technologies, Systems and Applications, July 2005

[4] Harris Interactive & SunGard Availability, “Survey of Fortune 1000 Companies Reveals Serious Deficiencies in Disaster Preparation - Troubling differences found between executives’

[5] The Conference Board, “Security in Mid-Market Companies: The View From The Top Executive Action”, 2004, www.conference-board.org

[6] Kibildis, George, “Business Continuity Planning in the Real World”, Disaster Recovery Journal, 2005

[7] Contingency Planning & Management / KPMG Business Continuity Planning Survey," cited in Andy Hagg, "BCP on the Rise," Contingency Planning and Management, January 2001

[8] Continuity Central, Vertitas, “Three-quarters of companies leave disaster recovery planning solely to the IT department”, www.continuitycentral.com/news0523.htm, 2005

[9] Mirifex Systems, LLC & Center for Regional Economic Issues Weatherhead School of Management Strongsville, Ohio , "An Analysis of the Consequences of the August 14th 2003 Power Outage and its Potential Impact on Business Strategy and Local Public Policy", 2004

[10] H. F. Lipson and D. A. Fisher, "Survivability-A New Technical and Business Perspective on Security," *Proceedings of the New Security Paradigms Workshop*, September 21-24, Association for Computing Machinery, 1999.

[11] Arnold, Richard L., CBCP, Prepare for Another Busy Hurricane Season:”, Disaster Recovery Journal, Volume 18, Number 3, 2005

[12] AT&T and the International Association of Emergency Managers (IAEM), “AT&T Study Finds U.S. Businesses Unprepared For Disaster” and “Disaster Planning in the Private Sector: A Look at the State of Business Continuity in the US”, http://www.att.com/news/2005/09/12-2, 2005

[13] “Designing Disaster Tolerant MC/ServiceGuard Clusters”, B6264-90002, June 1998, Hewlett-Packard Company

[14] F. Chang, M. Ji, S.T Leung, J. MackCormick, S. Perl, L. Zhang, “Myriad: Cost-effective Disaster Tolerance,” Proceedings of the FAST 2002 Conference on File and Storage Technologies. USENIX Association. Monterey, CA. January 2002.

[15] “Improving system availability with storage area networks,” Barocade Communications Systems, Incorporated, 2001. (http://www.dlt.com/storage/WhitePapers/Brocade/HA_WP_02.pdf)

[16] D. Fruend, “Disaster tolerant Unix: removing the last single point of failure,” Illuminata, Inc, 2002. Accessed at http://h71000.www7.hp.com/openvms/whitepapers/Illuminata.pdf

[17] K. Parris, “Disaster Tolerant Cluster Technology and Implementation”, HP World 2003 Solutions and Technology Conference and Expo, 2003.

[18] “HP Extended Cluster for RAC-100 kilometer separation becomes a reality,” A White Paper, Hewlett-Packard Development Company, L.P, 2004.

[19] Federal Emergency Management Agency. “Purpose of Standard Checklist Criteria For Business Recovery”, (no date). Retrieved February 12, 2005.