Objective

As availability requirements increase and Web-based applications, e-commerce, and enterprise resource planning (ERP) systems propagate; recovery timeframes must often reduce in geometric proportion. To compound the problem, many applications have evolved from being productivity aids to become absolute pre-requisites for the continuation of critical business processes. Today, recovering "critical" systems in 24, 48 or 72 hours is often an oxymoron as businesses are faced with constantly shortening recovery timeframes.

Issues

This heightened dependency imposes significant strain on the traditional shared-infrastructure model of recovery and forces contingency planners to find "new ways" to shorten recovery timeframes from days to hours...or even minutes!

By definition, the shorter the recovery timeframe, the more accessible the recovery hardware and software must be. Unfortunately, the more accessible the hardware infrastructure is to your company, the less it can be shared by others and the less effective a traditional, shared-cost hotsite solution is. Even many high availability solutions, like data shadowing or mirroring, only address a small portion of the problem. It's not enough to have data readily available if it still takes days to "connect" to your users and run batch processing. As a result, most businesses are faced with significantly higher infrastructure costs or are forced to compromise on longer recovery timeframes.

Solution

The solution to shortening recovery timeframes is conceptually the same regardless of whether you are running a mainframe or open system environment. In either environment, WTG shortens recovery timeframes by adjusting any or all of the three "Rs" in the formula: Reaction Time + Restoration Time + Recovery Time = Recovery Time Objective. The "new ways" of shortening recovery times are often process related and require that every possible minute of Reaction Time, Restoration Time and Recovery Time be reduced to the absolute minimum.

First, we focus on reducing or eliminating Reaction Time. Reaction Time is the time between the interruption and when recovery efforts begin. It consists of the time it takes to recognize the outage, assess the situation, mobilize personnel, identify the cause, evaluate possible solutions, declare a disaster, and begin recovery efforts. Reducing reaction time is a function of preemptive planning, understanding business impacts and commitment. Assuming an effective immediate reaction plan is in place, management must simply "pull the trigger" (e.g. declare the disaster and initiate recovery efforts) at the earliest possible time. Reaction Time is where some of the biggest times savings can come from...but to gain those savings often requires management to "pull the trigger" much sooner than they would typically be comfortable doing. Increasing their comfort level requires providing a complete understanding of the impact of not "pulling the trigger" and an implicit understanding in the very real cost of downtime. Our experience indicates that an IBPD (Iterative Business Process Decomposition) is the only effective way to understand the Recovery Time Objective (RTO) and to get management's buy-in. An IBPD clearly shows what happens to the business if interrelated processes are interrupted, and it defines impact in terms management understands as opposed to the arcane risk and financial impact estimates of a traditional Business Impact Analysis (BIA). Our IBPD also clearly quantifies inter-process dependencies and identifies which processes must be available for continuous business. This knowledge is critical when it comes to making the cost and capacity decisions necessary to address the high availability solutions required to achieve today's extremely shortened Recovery Time Objectives. If $1M is inarguably lost every 24 hours, it becomes much easier for management to immediately trigger recovery efforts, even if they should turn out to be preemptive.

Secondly, WTG designs a Data Availability Architecture to facilitate reducing or eliminating data Restoration Time. Restoration Time is a function of the amount of data, the frequency of backup, the effort of restoration, the amount of re-synchronization required and the accessibility to data. Depending on the platform, WTG evaluates a solution set of eight distinct data availability topographies to determine the most cost-effective based on business requirements as defined by the IBPD. These topographies include: periodic batch backup, electronic vaulting, remote journaling, real-time roll forward, standby database updating, symmetrical and asymmetrical replication, and synchronous and asynchronous mirroring. Each topography has its own pros and cons, and each offers a specific level of data availability. Higher levels of data availability, by definition, require shorter restoration times, which in turn serve to reduce the overall recovery timeframe.

Finally, WTG designs a customized System Availability Architecture that reduces Recovery Time. By crafting a physical architecture in consideration of hardware accessibility, software accessibility, switching speed and switching transparency, WTG is able to recommend and implement architectures ranging from HR (High Recoverability) to HA (High Availability) to CA (Continuous Availability). Our solution set addresses twelve distinct levels of increasing capability that range from recoverability through continuity and includes: cold sites, warm sites, hot sites, fault tolerant

configurations, redundant environments, clustering and hot fail-over, shared processing, standby data bases, transaction splitting, disaster resistant distributed processing, hot standby and full parallel processing. By crafting the right architecture in relation to the IBPD's defined critical processes and Recovery Time Objectives, recovery timeframes can be significantly shortened, production availability can be concurrently increased, and costs can be proportionate to the potential impact.

Scope

The strategies and techniques described in this solution are equally applicable to all levels of implementation, including: single business process, individual applications, single servers, platforms, whole sites or the entire enterprise.

Proven Results

WTG has employed the techniques described in this solution to help our clients implement recovery timeframes as short as several minutes. Recently, we have applied these techniques at a major auto manufacturer to design a multi-tier, high availability solution that reduced first-tier recovery time to less than ten minutes for all disasters short of complete plant destruction, and a second tier recovery time of less than 24 hours in the event of complete destruction of the physical plant.

WTG
THE
WILLIAM TRAVIS
G R O U P

The William Travis Group, Inc.
1827 Walden Office Square, S220
Schaumburg, Illinois 60173
847.303.0055 Fax 847.303.0378
www.williamtravisgroup.com

When it comes to highly-resilient, highly-available and highly-recoverable solutions for systems, data and people, WTG are the experts of choice. We design, implement and maintain cost-effective, leading-edge solutions within the framework of our holistic NextGen 360° ABC™ methodology so you can respond to any incident-any time. Our 360° NextGen ABC methodology can improve recovery and continuity planning results for "new players" or "old pros" alike and our continuity architectures are specifically designed to provide proportionate solutions that cost less.