

# Inside Perspective

Helping you unleash the full power of MEDITECH

Electronic newsletter published every 8 weeks

## Disaster Recovery, Business Continuity, and Reality

Jim Fitzgerald, Chief Technology Officer

Now that some parts of HIPAA are in force, and as we collectively continue to process and adapt to our 9/11-based emotional and mental scars, we may find ourselves asking the age-old question: "What if?" What if it becomes impossible to run the HCIS from the main data center due to bioterrorist activity or a government-mandated quarantine? What if a fire or flood knocks out my below-water-level data center? How will the nursing staff's workflow be impacted if electronic charting goes off-line?

We share your concern. Even better, we're doing something about it. Our *Disaster Recovery Working Group*, populated by a cross-section of our engineering, consulting, and technology management staff, is crafting a set of solutions that are as technically bulletproof as possible, yet flexible enough to adapt to different recovery objectives and environments. This article will introduce you to some of the concepts we've considered in our analysis and planning.

### *Disaster Recovery or Business Continuity?*

One of the first questions we asked ourselves was how to define the concept we were trying to develop. Was it *Disaster Recovery* or *Business Continuity*? Our director of engineering, Hugh Conway, put an end to the argument with the following observation: "Disaster Recovery is the process that leads to Business Continuity." And so it is.

### *Speaking the Lingo*

Early on in our discussions, it became clear that every organization will have different requirements, and that there needs to be a way to quantify them. If you are considering Business Continuity alternatives in your organization, it's important to understand several metrics, including *Recovery Point Objective* and *Recovery Time Objective*.

Recovery Point Objective (RPO) defines the maximum tolerable data loss resulting from a declared disaster situation, measured in hours or days. For example, if an organization is comfortable with restoring the prior evening's backups in the event of an emergency, their RPO is 24 hours. (On average, assuming a 6-hour backup window, the true RPO for this methodology will be in the range of 9 hours, but remember; we've chosen to define the maximum tolerable loss.) Organizations requiring the most disaster-tolerant solutions may decide that their RPO is zero—or effectively, no data loss.

Recovery Time Objective (RTO) defines the maximum tolerable delay to restore services. By default, RTO becomes a subset of RPO. Let's consider a tape-based recovery scenario with a 24-hour RPO. If the backup point is 12 am, and the known tape restore time is 6 hours, then a disaster declared after 6 pm will "break" the RPO, because the restore process will overrun the defined window for data freshness. Looking at our other example, if an organization is looking for an RPO of 0, they'd need to put together an infrastructure that has an RTO of 0, or they could not achieve their recovery goals.

### *The Illusion of Simplicity*

MEDITECH Magic customers live in a uniquely privileged world. It's a world where the operating system is lean, efficient, and complementary to the applications environment, a world where predictable operations are the norm. Because of these qualities, as well as the fact that network configurations are held almost entirely within the Magic OS, it is not uncommon for customers to decide that all they really need for Disaster Recovery and Business Continuity is a second set of servers in an alternate, readily-accessible location. While this can be true (assuming you're comfortable with a 24-hour RPO), the potential "rush to a conclusion" can often leave key questions unanswered: Is MEDITECH dependent on other interfaced applications? How will they fail-over? Will the network go down if the main data center is lost? How can traffic be re-routed? What about remote access users or sites connecting via WAN or VPN technology? How will IT staff operate systems in the event of a declared disaster? Will a fail-over site

affect departmental process or procedure? Clearly, there is more here than immediately meets the eye.

### *Today's Reality*

While it's possible to construct any number of scenarios—cold site, hot site, onsite, remote—which provide effective recovery of HCIS applications, Business Continuity can only occur where careful planning has created an environment where all the relevant infrastructural and organizational factors are aligned properly around a carefully thought-out recovery plan. If you are wrestling with these issues in your organization, I hope you'll talk to us. We're building a multidisciplinary approach to Business Continuity that we believe will be very effective, and our team is extremely interested in customer input. Please call, e-mail, or plan to visit us if you'd like to share your views.

*Jim Fitzgerald co-chairs our Disaster Recovery/Business Continuity Working Group. He can be reached at [editor@jjwild.com](mailto:editor@jjwild.com).*

Published by [JJWild](#)  
Copyright © 2003 JJWild. All rights reserved.