

Inside Perspective

Helping you unleash the full power of MEDITECH

JJWild

Electronic newsletter published every 8 weeks

Planning for the Unknown, Installment Three: Is Self-Hosted Disaster Recovery Right for Your Organization?

Sara Schaeffner, Senior Product Manager

Planning for the unknown is one of the most challenging things an IT Director or CIO has to do. It's hard enough to get funding for clinical equipment or 21st century computing infrastructure, let alone spend money on equipment or services you "may never need." But like it or not, the need for Healthcare Operational Continuity (HOC)—Disaster Recovery, Business Continuity, and High Availability—is today's reality.

In the previous two installments of this mini-series we suggested steps you can take to build a cost-justified case to invest in HOC, and to prioritize elements of HOC planning to arrive at an actual HOC solution that meets your organization's needs. In this article, we'll look at one of the distinct areas of HOC—Disaster Recovery—and focus on self-hosted disaster recovery solutions.

Disaster Recovery

In general, Disaster Recovery ("DR" here, but not to be confused with MEDITECH's Data Repository) is the process of recovering data and/or systems and resuming operations after an outage that was not prevented by some other element of an HOC strategy (like fault-tolerant systems).

We can all think of examples of DR scenarios: a flooded data center, major damage to IT systems and infrastructure caused by extreme weather or fire, or even a facility quarantine that prevents access to and continued production within the primary computing environment. These scenarios could also include something as simple but detrimental as database corruption that results in system failure and the need to restore to yesterday's data via tape—potentially wiping out up to 24 hours of productivity.

When you start designing a disaster recovery solution, one of the first things you have to decide is what level of responsibility the hospital itself wants to, and can, assume.

Basically, DR can be administered in the following ways:

1. The hospital can outsource to a managed service provider, e.g., JJWild's JSite.
2. The hospital can own responsibility for the disaster recovery process, e.g. self-hosted disaster recovery.
3. The hospital can have both managed *and* self-hosted DR, which many organizations feel is critical for protection against regional disasters.

Self-Hosted Disaster Recovery

Self-Hosted Disaster Recovery (SHDR) is not a one-size-fits-all "product" or even a

whole solution unto itself. Rather, it is an *approach* to disaster recovery, and is simply the opposite of outsourcing. Basically, SHDR is the do-it-yourself approach; if it's not hosted or managed by someone else, then it's self-hosted.

SHDR can take many forms and, by the above definition, could be as simple as a hospital purchasing warm spares or even just having backup tapes in an offsite location. However, these types of solutions might only address limited “disasters” and would not allow a hospital to achieve reasonable recovery times in the event of a major outage.

In fact, an SHDR solution that is designed to prepare an organization for mid-to-large scale events that could incapacitate the primary data center or a significant portion of the infrastructure within that data center should not leave facility preparations, equipment acquisition, or even just building and networking servers to the last minute. Most often, an SHDR solution that proactively addresses these elements will include a secondary data center (SDC) that is equipped to meet the environmental, network, and infrastructure recovery needs of the organization.

Due to the costs and staffing requirements associated with maintaining two data centers, some hospitals opt for outsourcing to a trusted managed service provider. But hospitals that do have an SDC (or approval for one) are half way to a robust SHDR solution!

Is there such a thing as “best practice”?

While SHDR can take many forms, some solutions meet the needs of the MEDITECH community better than others. And although “best practice” is a term that comes under constant fire, there are solutions that seem to meet these needs best. For example, many hospitals have opted to implement Storage Area Network (SAN)-based backup using Integrated Serverless Backup (ISB) and/or Integrated Disaster Recovery (IDR) technology for recovery at pre-built, pre-configured SDC facilities. Why?

ISB leverages the SAN to capture a “point-in-time” image of the MEDITECH database and to subsequently send this image off to tape or another storage medium—again, using the SAN to do the heavy lifting of backup, instead of the file servers themselves. This improves MEDITECH server performance during backup, and network performance in cases where hospitals are moving away from LAN-based backup schemas. It also means that backups can be performed once or even twice a day during “high traffic” periods without impacting users.

ISB provides other MEDITECH platform-specific benefits as well. Using ISB, MAGIC hospitals can, for the first time, implement an automated backup solution for MEDITECH, or even an enterprise backup solution that includes MEDITECH and non-MEDITECH applications. For Client/Server hospitals, the image-based, block level backup capability of ISB provides a greatly improved backup process for the Electronic Medical Record (EMR) which, due to the sheer magnitude of files, has historically struggled with file level backup technologies.

IDR takes the ISB technology a step further and gives customers the ability to create exact point-in-time copies of their MEDITECH systems on spinning disks that are bootable in the event of a disaster. This process can be scheduled several times a day, allowing hospitals to significantly reduce their Recovery Point Objective, meaning that the organization would lose less data in an outage, *and* their Recovery Time Objective since restoration from IDR virtually eliminates the need to restore from tape in a disaster.

Getting extra bang for your SDC buck

Interestingly, many hospitals that have a secondary data center for SHDR and are using technologies such as ISB/IDR or SAN mirroring are also seeking to utilize it for operational continuance purposes. Herein lies the distinction between two of the 'legs' in our HOC tripod: While DR is intended to facilitate recovery after an outage, operational (or business) continuance strategies are intended to maximize uptime by preventing total outages and/or by minimizing planned downtime unrelated to disasters.

Some customers that have an SDC for SHDR actually design their solution so that they can not only recover after a disaster, but also use the facilities and infrastructure to alternate between their primary and secondary data centers to perform upgrades and system maintenance without requiring prolonged downtimes. This type of solution can be ideal, but does introduce intricacy to the solution design and network topology that must be approached carefully and factored into all analyses and process development.

Beyond the data center—the people and the process

But, lest we forget in all our technical lust, the ability of an organization to recover from a disaster depends on more than just technology. You can spend millions of dollars on a robust SHDR solution in a phenomenal SDC but if no one knows what to do in a disaster, your users may never even know it was there.

Consider the sage advice and first-hand experience of Steve Sawyer, MEDITECH Utility IT Director for Mayo Health System, "A recovery plan that is not properly documented and also easy to use and access is basically worthless to any organization. These plans need to be tested and updated on a regular schedule by different staff members. Our recovery plans do not hinge on one person. We have a number of individuals who can and have executed the plan. This is possible because we have made the document and the exercising of that document a high priority in our work group."

Indeed, organizations should have multiple copies of clear How-To-Recover procedure documentation and these should be readily available to assigned resources and kept current. Organizations should also have test plans that are executed on a regular basis, and disaster recovery plans that include emergency contact information. On the other side of the workstation, each department should create and regularly validate its downtime procedures to ensure that users have instructions on how to perform their duties—including how to treat patients, administer medication, etc.—if mission-critical IT systems go down.

As Steve suggests, a disaster scenario is *not* the time for a trial run of your processes. IT and non-IT staff should be trained on how to recover mission-critical systems and restore application functionality as quickly as possible. And users need to be trained on how to function in an emergency. There's nothing like actually doing something to really learn it.

Self-hosted disaster recovery, managed disaster recovery, ISB, IDR, SAN mirroring, high availability, recovery procedures, test plans, business and network impact analyses, application criticality analyses, departmental downtime procedures—it's a big, wild HOC world out there.

But, starting with a thoughtful assessment of your organization and its real needs and capabilities, you can get to the right Healthcare Operational Continuity solution. And you really can plan for the unknown.

JJWild is here to help if you need us. Thanks for reading this series. I hope you've found it helpful, and would welcome your feedback.

Sara Schaeffner is Senior Product Manager for JJWild. Her product focus is on JJWild's Managed Service and Healthcare Operational Continuity offerings, including High Availability, Business Continuity, and Disaster Recovery solutions.