

Inside Perspective

Helping you unleash the full power of MEDITECH

JJWild

Electronic newsletter published every 8 weeks

Planning for the Unknown, Installment Two: Prioritizing Disaster Recovery Needs and Getting to the Right Decision for Your Organization

Sara Schaeffner, Senior Product Manager

In our first installment of this mini-series, “*Getting Your Organization to Get Serious About Disaster Recovery*,” we focused on steps you can take to build a cost-justified case for investing in reliable Healthcare Operational Continuity (HOC) strategies. (If you missed it, “[control click](#)” on this link:

http://newsletter.jjwild.com/e_article000639570.cfm?x=b7V7L1p,bRbD3Mp). But demonstrating that your organization is truly at risk—and that even non-catastrophic outages cost a lot of money—is only the first part of the process. You still need to come up with a specific plan.

This can be tricky. HOC strategies can be smaller than, as big as, or exponentially larger than your average breadbox. Fortunately, the data you gathered during your Risk Assessment and Business Impact Analysis (BIA) will also facilitate the next steps in the process: identifying your Recovery Time Objective (RTO) and Recovery Point Objective (RPO), defining the best options for your organization’s strategic HOC plan, and getting departmental participation and buy-in during the solution selection process.

First Comes First—Really, It Has To

The level of HOC fortification that you need will largely be driven by your organization’s dependency on IT systems. Before you can determine the true scope of your HOC requirements, you need to know the criticality of your systems and applications. Using the information collected during your BIA, perform an **Application Criticality Analysis** and prioritize your systems according to level of importance to the organization.

Categorize those systems that you absolutely cannot afford to lose as Tier 1; those systems that you require back in, say, two hours as Tier 2; within eight hours, Tier 3, and so on. Some of the criteria to consider when prioritizing:

- Is a system a single point of failure?
- Does it house clinically-sensitive applications?
- Is it part of an interdependent configuration?

Input from your departments will be extremely helpful in understanding what would happen if a system were to (hypothetically) go away. You may even want to solicit direct

inputs regarding the actual amount of time that users in a particular department could be without system functionality. (A word to the wise: When asked how long they can be without a system, users of any system to which they've become accustomed will inevitably say, "Not at all!" Time for some tough love. It's essential that you stress realism in your analysis. If all systems and applications end up as Tier 1, the price tag will likely kill the project altogether and, with it, the ability to protect your truly mission-critical systems. If you insist on a realistic approach, what you will ultimately get is a list of the systems that you *must* shore up in order to safeguard your organization.

How Long Is Too Long?

As you go through the Application Criticality Analysis exercise, you will simultaneously **define your organization's Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO)** for each data system.

An organization's RTO is the amount of time a system can remain offline without a significant impact to the business. It is calculated from the point of the outage to the point of recovery or resumption of operations. For example, let's take a hospital with a defined RTO of 12 hours for a particular system. If an outage occurred at 6:00 PM, users would need to be back on the system and functional by 6:00 AM to meet the organization's RTO.

RPO, on the other hand, represents how much data the hospital can lose without a significant impact to business and is calculated from the point of the last backup or data capture to the point of outage. Using our sample hospital and system above, let's assume they have an RPO of 24 hours and that they complete a nightly backup of that system at 4:00 AM every morning. Since the outage occurred at 6:00 PM, they would have lost 14 hours of data; this would be within the 24-hour RPO requirement. (It's important to note, however, that for many mission-critical applications a loss of 14 hours would, indeed, create a significant impact. For these, backing up once every 24 hours may very well not be enough.)

Your RTO and RPO will actually define key parameters of the technical and operational elements of your HOC strategy. So it is absolutely essential to have a clear understanding of these requirements.

Elements of True Healthcare Operational Continuity

Healthcare Operational Continuity is comprised of three parts: disaster recovery, business continuity, and availability solutions. Thus, a true HOC plan will consider:

- Production environment: data center caliber and fault tolerance
- Infrastructure: How current is it? What level of availability?
- Network: Accessibility and usability
- Data availability strategies: data backup or replication
- Disaster recovery facility and infrastructure: your own, a managed service, or both in the case of hospital-owned recovery centers that are geographically near to production facilities
- Disaster recovery plans and procedures

- Change management processes
- Departmental downtime procedures

In the last few years, MEDITECH hospitals have been introduced to several revolutionary HOC solutions, including:

- Integrated Serverless Backup (ISB) and Integrated Disaster Recovery (IDR) for data availability
- Stratus Technologies for server availability
- JJWild's JSite for MEDITECH-specific fully managed disaster recovery

Not all hospitals are able to go all-out right away, but it is possible to build a phased—often multi-year—strategic plan to get where you need to go. Keep in mind, too, that there's often more than one way to get there.

So, how are you to choose?

Sharing the Decision Process?

Decision-by-committee is not always practical, but in this case, you might find sharing the decision process to be a real help.

Consider the experience of Delnor-Community Hospital, a MEDITECH hospital in Geneva, Illinois. Hasi Smith, Director of Information Systems, and her team led a rigorous disaster recovery solution selection process that included other non-IS parts of the organization. The team developed three disaster recovery solution options and discussed how a disaster in each scenario would impact the hospital's business and reputation. They then engaged their super-users and department leaders *for each module* in a Maximum Allowable Downtime (MAD) exercise in which they simulated the impact of an outage during a critical time (for example, when working with the payroll manager, they simulated a disaster on a payroll Monday) and the department's ability to recover from it using each potential recovery solution.

Following the exercise, Hasi and her team compiled their MAD results and presented the options to the IT Strategic Council, including actual achievable recovery times, human resources required to recover, and cost for each solution. Equipped with this data, the Council was able to make an informed decision that significantly enhanced Delnor's Healthcare Operational Continuity readiness.

Delnor ultimately enacted a phased strategic plan that began with implementation of EMC SAN technology with BridgeHead Integrated Serverless Backup. The solution improved data availability and delivered improved backup times, performance, and RTO. For disaster recovery, they opted for JSite Managed Disaster Recovery Services. In the next phase of their plan, Delnor will adopt Integrated Disaster Recovery technology for both operational continuity locally and remote disaster recovery to the JSite facility—allowing for RPOs in the single digits.

Delnor's experience demonstrates that a key step in HOC planning and *execution* is to engage users in the process. By doing so, they become aware of the critical importance of disaster recovery and business continuity and the driving force behind it.

Your Plan to Plan for the Unknown

You've now got a roadmap to get you from where you are today—even if that's just simple tape backup—to true Healthcare Operational Continuity. To summarize, it looks like this:

1. Perform a risk assessment
2. Conduct a business impact analysis
3. Calculate the cost of downtime
4. Figure the incidence of downtime
5. Enlist organizational support (make sure clinical leaders are aware of your goals and the reasons behind them)
6. Perform an Application Criticality Analysis
7. Define your RTO and RPO for each data system and application
8. Design an HOC strategy that will meet these goals (including options)
9. Involve users in the final solution selection process
10. Arrive at a decision that is right for your organization and that will stick

In the third installment of "Planning for the unknown," we'll discuss an HOC approach we refer to as "self-hosted disaster recovery" (SHDR). A self-hosted disaster recovery solution—i.e., one that is owned and managed by the hospital organization as opposed to a managed service provider—can take many forms. It may actually accomplish much more than "just" disaster recovery, but it also must include the plans and procedures to make it work. Stay tuned for the next edition of *Inside Perspective*.

Sara Schaeffner, Senior Product Manager, is responsible for Managed Service and Healthcare Operational Continuity offerings, including High Availability, Business Continuity, and Disaster Recovery solutions. She can be reached at editor@jjwild.com.