

EDUCAUSE Center for Applied Research

Research Bulletin

Volume 2007, Issue 4

February 13, 2007

# IT Recovery Time Objectives: A Survey of Higher Education Practices

Ronald Yanosky, ECAR



In higher education information technology (IT) environments, what goes down *must* come back up. Thanks to lavish investments in enterprise systems, course management systems, portals, and other resources in the past 10 years, doing administrative or academic work today often is synonymous with using IT systems. As a string of recent events has shown—from the World Trade Center attacks to Hurricane Katrina—those systems may be even more vital to mobilizing an institution and keeping it going when disaster strikes. Yet IT systems, too, are subject to disruption. To a significant extent, the question of whether an institution can continue academic and business operations in an emergency is the question of how quickly key IT systems can be recovered.

Business continuity (BC) frameworks and standards commonly advise that enterprises answer those questions in part by assigning each key system a documented recovery time objective (RTO)—that is, a time period within which the enterprise plans to restore that system after suffering a disruption (ASIS International, 2005; ISO/IEC, 2005; Office of Government Commerce, 2001; Swanson, Wohl, Pope, Grance, Hash, & Thomas, 2002; IT Governance Institute, 2005; see also the draft BC management standard submitted to British Standards, 2006). Depending on the impact that an outage may have and the cost of architecting a recovery time frame, RTOs may vary from a fraction of a second to a period of months.<sup>1</sup>

RTOs are only one part of institutional business continuity planning, but they are a critical part because they distill much of the convoluted BC management process into explicit recovery goals. Among their benefits, RTOs

- are a practical expression of business and academic process priority;
- are a key input when designing recovery options and strategies;
- provide parameters for testing institutional BC readiness; and
- help institutions recognize and think through dependencies and priority conflicts.

In a comprehensive survey of higher education IT support for BC conducted in May 2005, the EDUCAUSE Center for Applied Research (ECAR) gathered responses from 340 institutions in the United States and Canada on a wide range of BC issues, including RTOs and related practices. In this research bulletin, we present our findings on RTOs to give a sense of how institutions are approaching them and what benefits they might bring. (The full BC study is scheduled for publication in 2007.)

While we hope our results prove helpful, we also caution against assuming that the RTO time frames presented here represent *de facto* best practices for higher education. The BC planning literature on RTOs stresses the importance of getting a deep understanding of the institution's particular needs, informed by business and academic priorities as executives and functional units outside IT define them. Institutions are likely to find that a realistic and balanced appraisal of system RTOs demands tradeoffs, a broad range of recovery time frames, and reconciling different constituents' views of what really constitutes top priority.

## Highlights of IT Recovery Time Objectives

Along with many other questions about BC planning and experience, our survey asked respondents to tell us whether their institution had a documented RTO for each of 17 separate infrastructure systems and applications. Some major business applications, such as human resources (HR) and finance, were broken into component pieces (for example, payroll and benefits administration for HR) so that we could distinguish business processes that might have different continuity priorities. We also asked whether the institution had a formal process for setting RTOs.

### Patterns in Recovery Time Objectives

Two findings quickly stood out in RTO results. First, for individual systems, the “no RTO” rate was consistently at or slightly over the halfway mark, ranging between 49 and 56 percent for every item. (See Table 1, which is sorted in descending order of combined 0–4 and 5–24 hour time-frame rates—that is, the percentage of institutions expecting to recover within one day. The most frequently cited RTO time frame for each item is highlighted.) This pattern held up overall as well: only 54 percent of respondents reported at least one documented RTO. “Don’t know” responses ranging from 2 to 8 percent for various systems probably mean that our RTO counts understate the actual rate, but not dramatically.

**Table 1. System RTO Status (All Respondents)**

System	0–4 hours	5–24 hours	1–2 days	3–6 days	7–14 days	More than 14 days	Don't know	Not applicable	No RTO
Campus Internet connection	17.2%	14.2%	7.5%	3.9%	2.4%	3.0%	2.4%	0.0%	49.4%
Institutional Web site	18.0%	12.3%	5.7%	2.7%	3.0%	0.6%	3.6%	1.8%	52.4%
Campus network	16.9%	13.3%	5.4%	5.1%	3.0%	3.6%	2.1%	0.0%	50.5%
E-mail	13.5%	12.0%	10.5%	5.1%	3.0%	3.0%	2.1%	0.0%	50.9%
Voice telephony	15.9%	9.3%	5.4%	4.2%	2.4%	0.6%	3.9%	7.5%	50.8%
Course management system	7.5%	8.1%	6.9%	9.6%	5.4%	2.1%	4.5%	3.6%	52.3%
Payroll	5.4%	8.1%	12.6%	7.2%	4.2%	1.2%	5.4%	6.6%	49.2%
Central finance/accounting	5.1%	8.4%	11.1%	8.1%	5.4%	1.5%	5.7%	3.3%	51.5%
Student records/registration	6.0%	6.0%	12.9%	7.8%	5.7%	0.9%	5.4%	3.9%	51.5%
Financial aid	5.1%	5.4%	12.0%	9.0%	4.8%	1.2%	6.0%	4.2%	52.4%
Admissions	4.8%	5.4%	12.0%	5.4%	6.0%	2.1%	6.3%	3.9%	53.9%
Student billing and payment	5.4%	4.8%	12.6%	7.8%	5.7%	1.5%	6.0%	4.2%	52.0%
Purchasing	2.7%	6.9%	10.8%	7.8%	5.4%	0.9%	7.2%	4.5%	53.6%
Library management system	4.2%	4.5%	5.8%	8.5%	4.5%	2.4%	7.6%	7.0%	55.5%
Benefits administration	3.9%	4.2%	7.2%	10.2%	5.1%	1.8%	7.5%	7.5%	52.7%
Recruiting	3.6%	3.3%	7.3%	7.3%	4.2%	3.9%	7.6%	6.6%	56.2%
Grants management	1.5%	3.0%	4.6%	6.7%	4.6%	4.6%	7.0%	14.3%	53.8%

Why do so many institutions operate without documented RTOs? Some may feel that their IT environments are simple enough to be recovered without explicitly prioritizing systems and documenting recovery windows ahead of time. Others may see the threat environment as so complex and uncertain that predetermined RTOs will be meaningless. And some institutions may see more value in spending resources on things like architectural resilience or staff training than in RTO creation.

The dearth of documented RTOs at many institutions, however, may well be due in part to constraint rather than strategic choice. Institutions with at least one documented RTO tended to agree more strongly that they had the necessary funding and staffing for central IT support of BC than those with no RTOs. The nature of institutional BC planning efforts seems to be a factor as well. Institutions with documented RTOs were more likely to report in-progress or completed formal institutional risk assessments and central IT BC or disaster recovery (DR) plans. For example, while 58 percent of those reporting no documented RTOs had an in-progress or completed central IT BC/DR plan, 81 percent of those reporting RTOs had one.


A second striking result in our RTO findings was the clear dichotomy between infrastructure items, which stood at the top of the list when measured by the percentage of respondents reporting 0–24 hour time frames, and applications, which tended toward longer time frames. Among all the infrastructure items, the most common time frame reported was 0–4 hours, while among all the applications, the most common time frames were either 1–2 or 3–6 days. (This pattern may be seen clearly by examining the shaded cells in Table 1, which identify the most commonly reported time frame for each item.)

The contrast between infrastructure and applications may be seen in Figure 1, which shows the distribution of reported time frames among respondents reporting RTOs. (Because Figure 1 excludes “no RTO” respondents, the values for each time frame are larger than those reported in Table 1, and the sort order is slightly different.) As in Table 1, infrastructure items rise to the top of the 0–24 hour RTO time-frame rates.

Setting more aggressive objectives for infrastructure items makes sense. Authorities on emergency preparedness and BC stress the vital importance of communication during a crisis, and the campus network and the Web and e-mail traffic it carries are a communications resource of the first importance. Availability of these items is also often a precondition for the recovery of enterprise systems. Since many institutional emergency plans call for staff to work from home, live network connections may actually be more important to continuity during a disruption than in normal circumstances.

More surprising than the priority of infrastructure over applications is the relative weakness of differentiation among the applications, and consequently a general similarity in their RTO profiles that may be seen in Figure 1. Payroll, course management systems, and central finance/accounting stood highest among applications ranked by 0–24 hour time frames, while recruiting and grants management sat at the bottom. But the time-frame rates for applications often varied within fairly narrow ranges. For 9 out of 12 applications, for example, the number of respondents citing RTOs under 24 hours was between 25 and 35 percent.

**Figure 1. System RTO Time Frames (Respondents Reporting at Least One RTO)**



This does not mean that there are no differences worth noting in the applications RTOs. It says a lot about rising academic dependence on IT that a relative newcomer among enterprise applications, course management systems, had a 24-hour time-frame rate slightly greater than payroll, which is often cited as a top priority in BC circles. But a

crucial companion to the course management system, the library management system, was considerably farther down the list. And the other key academic application we asked about, grants management, had the lowest 24-hour time-frame rate of all. Surprisingly, we did not find significant differences among Carnegie classes in grants management RTO time frames.

Of course, individual institutions may have much more graduated RTO profiles. It may be, too, that at some institutions, integrated ERP systems effectively allow (or require) that multiple business processes and applications be restored together, promoting a less granular approach to recovery. But the general absence of strong differentiation could also indicate that procedures for determining RTOs are not as robust as they might be and therefore aren't forcing difficult choices—or uncovering hidden priority conflicts—that would push some applications to longer time frames. A certain flavor of wishful thinking in reported RTO time frames adds weight to this speculation. In every application except grants management, two-day or shorter time frames accounted for 47 percent or more of RTOs, and six-day or shorter time frames accounted for 70 percent or more. Considering that some authorities warn that it may take several days to restore enterprise systems from backup media even to a well-prepared hot site—and only one in five of our respondents reported an operational hot site—these RTOs seem aggressive. The danger of falling into wishful thinking is one reason why RTOs, like all BC plan elements, should be subjected to regular and realistic testing.

### Formal Process and RTO Count


Having documented RTOs for systems doesn't necessarily mean having formal processes for creating them. BC standards and frameworks recommend that IT collaborate with functional departments when setting RTOs to ensure that the RTOs reflect a common understanding of business process priorities and that expectations are set appropriately. Typically, the RTOs might grow out of a business impact analysis in which relevant departments rank the priority of their systems using a standard scale. Collated, compared, and assessed in light of IT realities, these evaluations can become the basis for a set of RTOs, which might then be periodically reevaluated.

It appears, however, that where RTOs are being created, it's about as likely to be in the absence as in the presence of a formal process. Overall, only 25 percent of our respondents reported having a formal RTO process. Among institutions that had reported at least one documented RTO, 47 percent also reported having a formal process, and 53 percent said they didn't.

### Who Has RTOs?

In general, we found that bigger institutions and those offering more advanced degrees were more likely to report setting at least one RTO. The effect was strongest for institution size: while 32 percent of institutions of 1–2,000 students reported at least one RTO, 76 percent of those with more than 15,000 students did so (see Figure 2).

**Figure 2. Institutions Reporting at Least One RTO, by Institution Size**



Among Carnegie classes, about two-thirds of doctorals (68 percent) had at least one RTO, a bit short of double the rate (37 percent) among baccalaureate institutions. Master's and associate's institutions fell in between.

Within the total respondent population, we also found a tendency for larger institutions to report more RTOs. As Table 2 shows, the mean and median number<sup>2</sup> of RTOs reported grew along with institution size categories, from a mean of 3.72 (median 0) RTOs for institutions of 1–2,000 students to a mean of 9.71 (median 11) among those of more than 15,000 students.

**Table 2. Mean Number of RTOs by Institution Size**

FTE Institution Size	All institutions			Has at Least One RTO			Has Formal RTO Process		
	Mean	Std. Deviation	Median	Mean	Std. Deviation	Median	Mean	Std. Deviation	Median
1–2,000	3.72	6.248	0.0	11.82	5.25	15.0	13.43	4.353	15.0
2,001–4,000	5.54	7.036	0.0	11.42	5.86	14.0	12.92	6.388	16.5
4,001–8,000	6.20	7.332	0.0	13.15	4.65	15.0	14.67	4.451	16.0
8,001–15,000	7.33	7.302	5.0	12.00	5.53	14.0	14.00	5.237	17.0
More than 15,000	9.71	6.734	11.0	12.82	4.42	13.5	13.11	4.211	15.0
Total	6.57	7.188	2.5	12.34	5.07	15.0	13.58	4.778	16.0

But this difference was largely a result of the greater propensity of smaller institutions to report zero RTOs—one reason for the high standard deviations associated with these size-range means. When zero-RTO respondents were excluded, leaving only those institutions reporting at least one RTO, the mean count of RTOs ran in the range from 11.42 to 13.15, and significant differences between size ranges disappeared. Restricting the respondent base to those reporting a formal process for setting RTOs had a similar effect, except that the mean and median RTO counts were slightly higher. And when we did a combined comparison of these two factors—examining all those with at least one RTO by whether or not they had a formal RTO process—we found that those who reported RTOs but no formal process had a mean of 10.73, while those with both had a mean of 13.45.

In short, while our larger respondent institutions were more likely to establish RTOs and to have a formal RTO process, they did not tend to produce more RTOs than smaller institutions that also did those things, at least within our list of systems. Perhaps there is something about larger institutions—greater resources, more complex IT environments—that impels them more often to start creating RTOs. But smaller institutions often seem to set just as many (or more) RTOs as their bigger cousins when they undertake the task. If the goal is to have RTOs for most or all major systems, it may help to have a formal RTO-setting process. But most of all, it appears, it’s important to get started. RTOs seem to beget RTOs.

### RTOs and Perceived Preparedness to Restore

While our study captured the reported quantity of RTOs at respondent institutions, it couldn’t directly assess their quality or utility. Is there any reason to believe that RTOs contribute to the ultimate goal of BC planning—greater readiness to respond to a crisis?

We cannot give a definitive answer, but we did find some evidence that RTOs may contribute to the perception of being better prepared. Respondents who reported at least one RTO tended to agree more strongly, compared to those without RTOs, that their institution was prepared to restore centrally controlled systems in the event of a disruption (see Table 3). On a scale of 1 to 5, where 1 = strongly disagree and 5 = strongly agree, those with RTOs rated themselves on average about midway between a neutral (= 3) answer and agreement, while those without them averaged a straight neutral response. Having a formal RTO process combined with at least one RTO averaged a still higher 3.74 response.

**Table 3. Institution Is Prepared to Restore Centrally Controlled Systems, by RTO Status**

RTO Status	Mean	Std. Deviation
No RTOs	3.01	1.010
At least one RTO	3.49	1.010
At least one RTO, no formal RTO process	3.20	1.111
At least one RTO, has formal RTO process	3.74	0.833

*1 = strongly disagree, 2 = disagree, 3 = neutral, 4 = agree, 5 = strongly agree*

## What It Means to Higher Education

Our results suggest that IT administrators see value in documented RTOs. Institutions that set RTOs tend, on average, to set more than a few, and those that formalize the process set more. Though many variables influence respondents' overall sense of preparedness to restore centrally controlled systems, documenting RTOs probably plays a role in the substantial difference in preparedness self-ratings between institutions with and those without RTOs.

The aggregate pattern of reported RTO time frames is also instructive: judged by the assignment of short RTO time frames, respondents give highest priority to those IT items—Web site, telephony, network, and e-mail—that help the institution organize a crisis response and maintain or restore operational continuity. Though application RTO time frames are a bit longer and priorities less sharply defined, they seem arranged above all to mobilize staff and students and keep them productive, for example by putting payroll and course management systems high in the recovery order.

Challenging though it may be, it appears that the RTO process is worth the trouble—and so our most worrisome finding is that nearly half of respondent institutions reported no documented RTOs. Though more concentrated among smaller institutions, these respondents were present in every size range, and we found signs that lack of RTOs could in many cases be a result of a general underfunding of BC support in central IT. Presumably, institutions lacking documented RTOs have some layers of defense, whether in the form of undocumented RTOs, or staff skills, or robust architectures that foster smooth recovery. But even where the technical recovery process is strong, institutions lacking RTOs run the risks of a lack of shared understanding about recovery times between central IT, its users, and executive leadership. RTOs, in other words, can help manage expectations—a key to the *perception* of success.

After all, when a failure occurs, recovery times will emerge, whether planned for or not. Institutions that work with user departments to plan and document their RTOs—and then test them—are less likely to get a nasty surprise when they discover how long it really takes to recover crucial systems.

## Key Questions to Ask

- Which academic and business processes are the most important to maintaining constituent services during a disruption in operations?
- How can clashing perceptions about the priority of academic and business processes be reconciled?
- How does the presence/absence of comprehensive BC planning at our institution affect our ability to define and deliver RTOs?

- How do the institution's continuity objectives affect IT infrastructure, staffing, and planning?
- What tests and exercises will help us understand how realistic our RTOs are?

## Where to Learn More

- Cramer, B. (2004, May 4). Business continuity metrics: How much can you afford to lose? *Computerworld*. Retrieved October 18, 2006, from <http://www.computerworld.com/printthis/2004/0,4814,92865,00.html>
- Hiles, A. (2004). *Business continuity: Best practices*. Brookfield, CN: Rothstein Associates, Inc.
- Sun Microsystems. (2006). *Linking disaster recovery time objectives to business requirements*. Santa Clara, CA: Sun Microsystems, Inc. Retrieved October 18, 2006, from [http://www.sun.com/storagetek/white-papers/linking\\_disaster\\_recovery.pdf](http://www.sun.com/storagetek/white-papers/linking_disaster_recovery.pdf)

## Endnotes

1. A related concept beyond the scope of this research bulletin is the recovery point objective (RPO), the maximum acceptable interval between a backup and a potential interruption, during which data will be lost. When a system is recovered, the RPO defines how old the restored data may be relative to the time of the disruption.
2. The mean is computed by dividing the sum of a set of items by the number of items. The median is a value in an ordered set of values below and above which there is an equal number of values. If there is no single middle number, the median is the arithmetic mean of the two middle values.

## References

- ASIS International. (2005). *Business continuity guideline: A practical approach for emergency preparedness, crisis management, and disaster recovery*. Arlington, VA: Author. Retrieved October 18, 2006, from <http://www.asisonline.org/guidelines/guidelinesbc.pdf>
- British Standards. (2006). DPC BS 25999-1. *Code of practice for business continuity management*. London: Author. Retrieved October 18, 2006, from <http://www.survive.com/resources/displayresource.cfm?SubsubjectID=1&ownerID=470&subjectID=30&viewby=3>
- ISO/IEC. (2005). *ISO/IEC 17799 Second edition 2005-06-15: Information technology—Security techniques—Code of practice for information security management*. ISO/IEC. Available from <http://www.iso.org/iso/en/CatalogueDetailPage.CatalogueDetail?CSNUMBER=39612>

- IT Governance Institute. (2005). *CobiT 4.0 Control objectives, management guidelines, maturity models*. Rolling Meadows, IL: Author. Retrieved October 18, 2006, from <http://www.isaca.org/Template.cfm?Section=COBIT6&Template=/TaggedPage/TaggedPageDisplay.cfm&TPLID=55&ContentID=7981>
- Office of Government Commerce. (2001). *ITIL: Managing IT services*. London: The Stationery Office.
- Swanson, M., Wohl, A., Pope, L., Grance, T., Hash, J., & Thomas, R. (2002). *Contingency planning guide for information technology systems: Recommendations of the National Institute of Standards and Technology*. NIST Special Publication 800-34. Washington, DC: US Government Printing Office. Retrieved October 18, 2006, from <http://csrc.nist.gov/publications/nistpubs/800-34/sp800-34.pdf>

## About the Author

*Ronald Yanosky (ryanosky@educause.edu) is a research fellow at the EDUCAUSE Center for Applied Research.*

Copyright 2007 EDUCAUSE and Ronald Yanosky. All rights reserved. This ECAR research bulletin is proprietary and intended for use only by subscribers. Reproduction, or distribution of ECAR research bulletins to those not formally affiliated with the subscribing organization, is strictly prohibited unless prior permission is granted by EDUCAUSE and the author.