Incentives, Behavior, and Operational Risk Management in IT

Incentives are certainly not the be-all-and-end-all of operational risk management in IT departments. But most practitioners would agree that, when incentives don't work well—or work perversely—the chances of operational risk mishaps in IT increase. Moreover, IT has become a major part of financial services, and that translates into a good deal of risk exposure with potentially disastrous business consequences.

BY CHARLES TAYLOR

WHAT ARE THE main types of operational risk in IT departments? At the highest level, they can be described like any other operational risks in a financial institution: as the direct or indirect risk of loss from failures in people, processes, or systems or from external events. But because the nature of IT is to either build or operate systems, the risks can be described more precisely in terms of system design, build-and-deploy activities, and operation and maintenance.

The types of things that can go wrong are delays, cost overruns, failure to meet business requirements, and complex systems that simply cannot be operated at low risk. Add to those the risks involving failure to achieve required standards of availability, integrity, security, resilience (including business recovery), and adaptability in the face of change or mishap.

Whom to Incent?

While all IT people should be aware of the operational risks they take and be motivated to keep them within well-defined and reasonable bounds, it is not clear that everyone should receive financial incentives. Institutions that have begun to pay incentives for IT operational risk management have generally confined them to the department's senior management. There is a case, however, for broadening incentive pay to all key professionals, including systems analysts and architects, whose actions directly impact operational risks. Their incentives should, of course, be focused on the risks they affect—primarily investment-cycle-related risks, in the case of systems analysts, for example. Accordingly, their incentives should be structured differently from those for senior managers, but they also might be just as large, given that the professional misjudgment of analysts and others could have an effect on operational risk for years to come.

Are team rewards more just and effective than individual rewards? Team incentives encourage stronger members to cover for weaker ones, which can be positive for the culture and the organization in the short run, but may also give the "stars" an incentive to leave and low performers the opportunity to coast—unless the stars are celebrated and low performers managed aggressively. To work in the medium term, team incentives may well require stronger management skills.

This issue of whom to incent is part of a larger issue of how to establish a strong risk management culture in IT. Organizations approach this issue in different ways. Some get commitment and leadership from the CEO and the board, the most effective sponsorship possible. Some make risk management extremely visible.[1] And some

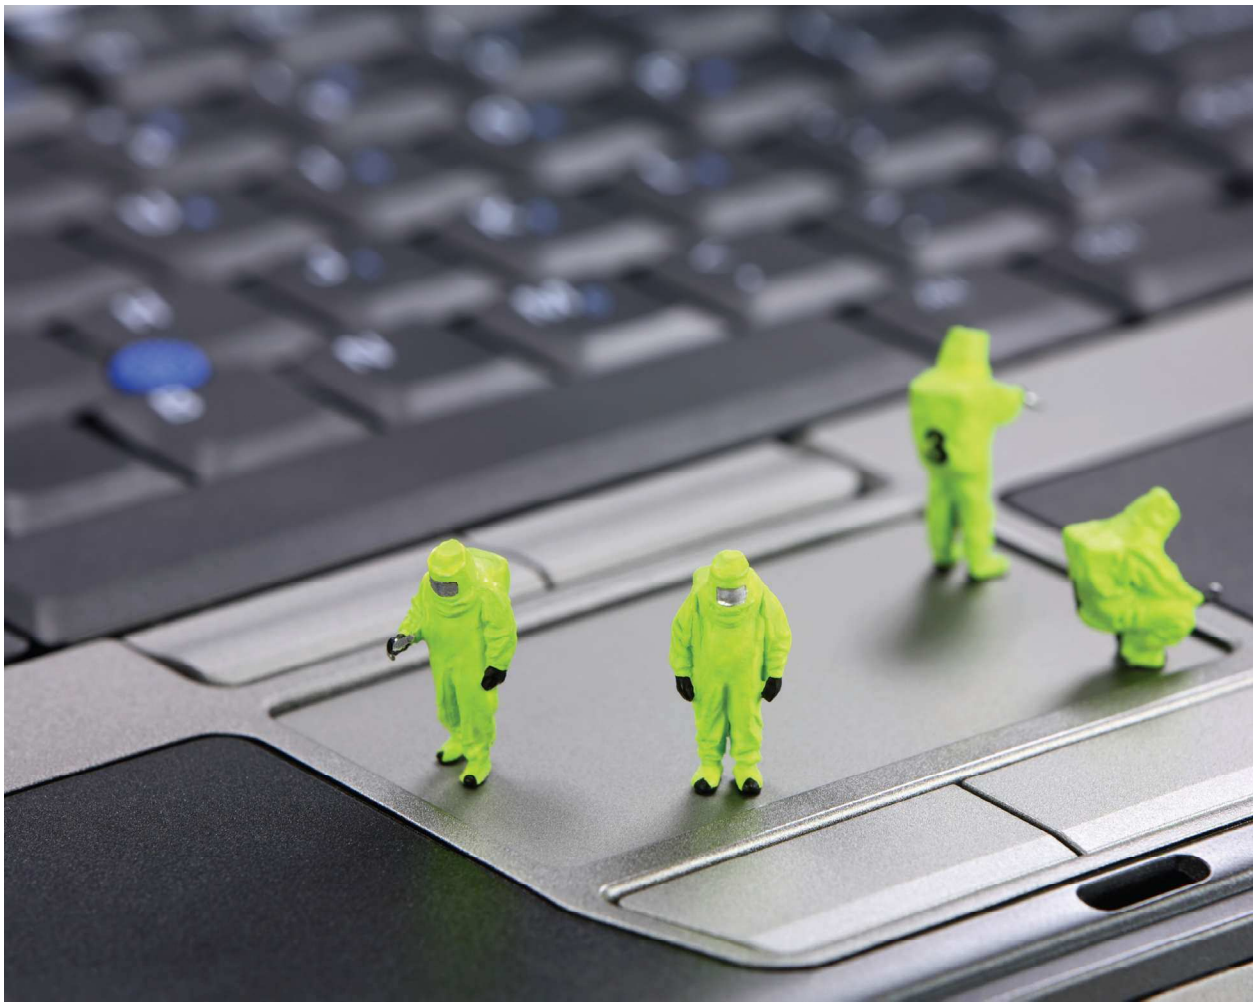

encourage identification, discussion, and management of risks—the kind of transparency that leads to systemic improvements.

Others, however, do not give operational risk management the same weight. They may reward the firefighters—the ones who make a valiant effort to deal with an unexpected event—but are unlikely to reward those who keep things on an even keel. Setting the tone at the top may well count as much as financial incentives in encouraging IT individuals and teams to manage operational risks well.

Drivers of Operational Risk in IT

The main drivers of operational risk in IT are complexity, the pace of change, and business volatility. *Complexity* is a problem often tackled by adopting a flexible underlying architecture for integration and application development. The problem is that creating and applying such an architecture becomes an increasingly delicate task as the business of a financial institution matures and the reach of technology increases. The pace of change in IT results from a combination of changing business requirements and changing underlying technology. Institutions in the throes of significant business change, such as a merger, will have higher-than-average IT risk. And ongoing *volatility* of business needs—for example, in the brokerage business, which tends to be much more volatile than, say, retail banking—creates more fluid and challenging demands on IT as new products are introduced and activity waxes and wanes over the business cycle.

Then there are some longer-term drivers of IT operational risk. For example, *aging* systems and infrastructure are an ever-present source of risk. As IT becomes more ubiquitous and as the discipline ages, so too does the stuff of which systems and infrastructure is made. Of course, institutions and their IT departments have programs to upgrade and replace much of it, but that does not alter the fact that the task of managing aging systems is a growing one and contributes to the accumulation of operational risk in IT.

Another longer-term driver of operational risk is the increasingly *virtual* nature of IT. The more virtual the environment, the more difficult it will be to detect certain kinds of problems. For example, in the past, stolen data

meant stolen files or paper—something easily seen and comprehended. Today, stolen data means that data has been copied, which may look just the same as data that has not been copied. Detecting theft may be intrinsically harder to do.

Moreover, increasing exploitation of the Internet means more *connectedness* between business systems and unprotected environments, especially if significant numbers of employees need and have remote access. In addition, interconnected computers offer malicious hackers the possibility of speed and economies of scale and scope. They can execute an attack in very little time across a wide range of systems and locations.

Still, at least one pervasive trend tends to reduce operational risk. The Internet in particular and greater *interconnectedness* in general create whole new levels of *redundancy* and *resiliency*. For example, voiceover IP over wifi may be an effective protection against landline problems; and instead of just losing hardware during a natural disaster, we can today manage this problem with "hardware extraction." Virtual machines are, in many respects, more stable than physical ones.

Alignment

All too often, IT professionals view their department as if it were divorced from the businesses it supports. That's not the way it should be. The IT department's risk culture and compensation should mirror the businesses of which it is a part. In a stable and controlled environment, risk can be calibrated precisely and managed effectively using stable IT measures that track critical business outcomes selectively and directly. For example, IT departments running large customer access networks might be well aligned if they were incented to ensure network availability, particularly in times of predictably high usage. But where business results depend on the flexible management of change, growth, and volatility, IT rewards for risk management (and performance) should be directly related to business results such as revenue and profit growth. The IT department supporting a volatile business should be incented to manage the risks around rapid development and deployment of new systems and around flexible capacity management.

Risk objectives should be related to business objectives as closely as possible. For instance, an outage indicator for a system that has an operating cycle calling for intense processing during a limited period each day should be measured during that period rather than averaged for the day as a whole. Likewise, outage indicators for critical systems ought to be measured separately rather than averaged in with other systems of different business criticality.

Finally, alignment is not a static thing. As business priorities change, risk priorities and therefore incentives need to change too. And an element of the incentive scheme should reward flexibility and responsiveness to unexpected shocks. That includes planning for and coping with disasters.

Risk Optimization versus Elimination

Whatever the alignment of risk objectives to business objectives, there is one other conceptual problem many firms face when setting incentives for operational risk management in IT. The problem is getting managers to strive toward optimizing the management of risks as opposed to eliminating them. For example, it is common to articulate an availability standard (or a risk of non-availability standard) as a percentage such as five nines, or 99.999%. Does that mean that a manager who achieves at least this level is managing risk well?

Not necessarily. First, the costs of achieving, say, six nines may be far too high—the desired range may be 99.999% to 99.9995%. Second, firms may not have any indicators that tell them how close they are to the edges of the desired range—how close they are to incurring too little as well as too much risk—at any point in time. Third, even if they do, they may not be managing to those indicators.

Most IT risks ought to be managed quite tightly to achieve a particular balance that represents the firm's chosen trade-off between risk and return. These risks should not, as a rule, be managed down to zero.

Risk Indicators

Incentives for managing operational risks have to be based on something measurable that management and key personnel can influence. Such yardsticks of behavior are usually referred to as *risk indicators*.

Some risk indicators are back-to-front performance indicators. As noted in the example above, keeping the risk of unavailability below 0.001% is the same thing as achieving availability of 99.999%. In other words, the unavailability indicator is nothing more than 1—the availability indicator. When it comes to how incentives incorporate these sorts of indicators, it doesn't really matter whether they are built into the incentive system as a performance measure or a risk measure as long as they receive their due weight.

But these upside-down performance indicators are not the whole story. For one thing, the volatility of the performance indicator is often a useful additional risk indicator. For example, how much actual availability *fluctuates* between, say, 99.999% and 99.9995% provides a useful sense of how *tightly* the risk around availability (or unavailability) is managed.

Other important risk indicators for incentive design have little to do with performance indicators. For example, indicators that are based on some count of virus attacks repelled and the depth of the defense penetrated are much more directly about disasters avoided than business performance enhanced. Equally, indicators around ID

IT incentives ought to reward good outcomes rather than simply the avoidance of bad outcomes. For example, IT should be rewarded for its contribution to high rates of customer satisfaction in telephone banking or in ATM usage. The person who prevents fires ought to get as much recognition as the person who puts them out.

management are entirely about control rigor. Key person attrition rates or even simple staff turnover rates might also be thought of principally as risk indicators in that they are associated with the risk of deterioration in service levels rather than the actual incidence of worsened service.

Whatever the kinds of indicators, they are likely to have thresholds associated with them, and these are often thought of as the boundaries between green, yellow, and red ranges, where green is good, yellow is so-so, and red is bad. These ranges need not be the same for all purposes. For example, for management and reporting, a threshold may mark a point at which an indicator is reported upward or the management of the underlying risk is escalated. For incentives, however, yellow might represent expected outcomes with little or no associated incentive reward, and green may be exceptional outcomes.

Consider the delivery of projects under unreasonable time pressure. It is not uncommon for CEOs or others in senior management to set deadlines for application development teams that can be met only with Herculean efforts. In these circumstances, is overcoming the risk of delay what is expected—and therefore not meriting incentive and reward? Or is it exceptional—and therefore worth financial incentives?

It really does depend on the prevailing culture of the institution and its IT department. If this is how business is always done, then it is hard to say the situation is exceptional and suitable for significant reward. Still, there is plenty of evidence that operational risks increase when people work under excessive pressure. So it seems likely that IT departments that are always in crisis mode will make more mistakes than average and find it next to impossible to sustain exceptional results.

Accordingly, senior management should select indicators and indicator thresholds for measuring IT operational risk management performance in conjunction with IT and business management. The thresholds should mark attainable performance standards that reflect the risk appetite of the institution as a whole. Involving business management in the selection and review of these indices should improve the chances of effective alignment and reduce the risk of unintended adverse incentives or of gaming.

Finally, standards of risk management and performance should generally rise over time, reflecting improvements in technology and management practices. Going and staying "green" does not necessarily constitute success if standards are static.

Structuring Incentives

Incentives and the indicators on which they are based should be as comprehensive and relevant as possible. They should reward the good management of all material IT risks, including delays, cost overruns, and failure to achieve required quality along with standards of availability, integrity, security, resilience, and adaptability. It matters little if some of these incentives are understood to be performance measures as long as all material risks are covered and each is accorded a degree of relative importance that aligns it to the objectives of the businesses being supported.

While incentives ought to be comprehensive, the weight given to different indicators in crafting the overall incentive scheme ought to reflect where risks are really concentrated.

Incentives for good risk management have to be material to incent behavior. Managers and key personnel may well ignore small incentives that have a minimal impact on compensation. But if some improvements in risk management are individually small but collectively significant enough to enhance performance, it may be worth including them in the incentive scheme.

Indeed, this issue may arise for operational risk as a whole. Its contribution to an overall incentive package may be so small that an aggressive manager might, for example, decide to give it next to no attention and focus entirely on getting projects completed on time and within budget.

Encouraging Proactive Risk Management

IT incentives ought to reward good outcomes rather than simply the avoidance of bad outcomes. For example, IT should be rewarded for its contribution to high rates of customer satisfaction in telephone banking or in ATM usage. The person who prevents fires ought to get as much recognition as the person who puts them out.

How can indicators and incentives actually reward people for things that didn't happen? Sometimes, as in the case of virus attacks repelled, there are objective measures of how serious a threat was, and these measures can be used to gauge the quality of risk management. Sometimes it may be possible to assess performance objectively against peers. For example, if yours is the only bank in the region to have successfully repelled hacker attacks over the past 12 months, that might be an indicator of superior risk management.

Sometimes indicators that have very little predictive value in isolation may reveal predictive trends—cycle time to recovery, security incident response times, or service issue resolution statistics, for example. In such cases, the quality of proactive management can be judged by how early a threatening trend was identified and how effectively it was thwarted.

Capital for Operational Risk in IT

Some institutions assign operational risk capital to the IT function, and this adds an additional incentive for senior IT managers to manage their operational risks well. This is difficult to implement, however. Today, it is more common for that risk to be covered explicitly or implicitly by capital assigned to the businesses IT supports.

Wherever it is cost effective to estimate the future distribution of loss associated directly with operational risks in IT, it is possible to calculate a capital charge for IT operational risks. If transfer pricing is used in an institution, then it is even possible to calculate a risk-adjusted return to capital in IT. But even if it is not, it might still be useful to calculate economic capital as a quantity to be managed down by IT and something to be allocated to supported businesses in their RAROC estimates.

Conclusion

Incentives can make an important contribution to how well IT operational risks are managed. But first, institutions have to identify their operational risks in IT, determine how material each one is to the businesses IT supports, figure out whom to incent (just IT management or all key IT professionals—and in teams or as individuals?), and come up with a plan for doing so (use thresholds for indicators imbedded in a balanced scorecard, or use estimates of economic capital?).

However it is done, systematic attention to the issue of incentives in IT is important. Increasingly, financial institutions will have to manage IT risk not only to operate efficiently and competitively, but also to ensure that their reputation for security, integrity, and service is nurtured and preserved in the marketplace. ❖

Charles Taylor is director of Operational Risk at RMA. Contact him by e-mail at ctaylor@rmahq.org.

Notes

1. For example, one organization represented in RMA's IT Working Group Forum considers customer-facing systems availability so critical to its Internet businesses that it has screens monitoring it displayed in its lobbies.