NOTICEBORED

Management briefing on

Contingency metrics

"A few well-chosen metrics can be a huge help in monitoring controls and measuring their effectiveness”

Clint Kreitner, SANS

Summary

This is a management discussion paper outlining potential metrics for measuring and improving contingency planning, resilience, disaster recovery and so on.

Introduction

Measuring the effectiveness of contingency arrangements is a tough challenge, not least because (like insurance policies) we hope we will never need to use them. However it makes sense to measure our investment in contingency plans and preparations, and to confirm whether management is sufficiently confident in them, prior to enacting them as by that stage it will be too late.

Contingency targets

Two parameters are key for contingency planning purposes as they drive important architectural and investment decisions:

- **RTO** - Recovery Time Objectives define the maximum acceptable period of interruption of business processes and the associated IT systems. Processes with very low RTOs have to be engineered for high availability and resilience, with full duplication or at least automated IT failover arrangements to minimize disruption if primary systems fail.
- **RPO** – Recovery Point Objectives define the maximum acceptable loss of live data. Processes with very short RPOs need data mirroring and similar near-real-time duplication of data between primary and secondary IT systems.

Costs increase dramatically as RTO and RPO values fall, hence the resilience and recovery objectives need to be matched against the potential business losses while systems/data are unavailable. In short, it makes no sense to spend more on the design, engineering and maintenance needed to support, say, a one hour maximum down-time than the business would probably lose in an hour under worst-case conditions.

There can be further hidden contingency costs for business- or safety-critical systems, namely vital parts of the supporting IT infrastructure, particularly the networks. Most modern networks are designed and operated for high availability by default but any single points of failure (such as particular firewalls, routers or WAN connections) are availability threats. Uprating them can be very costly but cannot be neglected.

Contingency metrics

Resilience metrics

Resilience, the capability to keep business processes running despite incidents, can be measured by identifying the rate and severity of unplanned interruptions and IT service outages that actually occur. Potential metrics include:

- The number of incidents that occurred during the reporting period, normally classified by severity and perhaps with trends (*e.g.* "We suffered two Severity 1 and twelve Severity 2 incidents during June, compared with one and three respectively during May");
- Some assessment of the IT and business costs incurred to analyze, stop and recover from incidents, including both tangible and intangible losses incurred (*e.g.* "The incidents in June are estimated to have cost $X in lost production, $Y in service penalties and $Z in recovery activities, representing a cumulative total loss for the calendar year to date of $$$"). Even if only estimated, these numbers provide a better basis for the investment decisions around contingency, resilience and recovery.

Case studies on particularly serious or unusual incidents serve to illustrate weaknesses or limitations in the controls and form an effective security awareness-raising mechanism in themselves. If well written, they can also demonstrate the value of contingency measures and reinforce the heroic efforts of employees to restore normal service.

Recovery metrics

The proportion of business processes and IT systems for which RTO and RPO have been defined and agreed is a crude guide to the maturity of contingency planning.

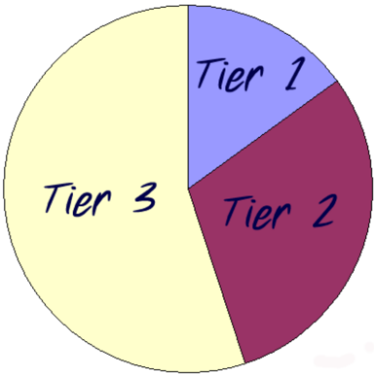

More detail can be obtained simply by comparing the numbers of processes/systems in "tiers" of availability – typically three categories are used with "tier 1" being the most critical processes/systems that require the highest levels of availability and resilience with the lowest possible recovery times. Broadly speaking, one would expect the proportions to be something like those shown on the pie chart (there is no formal theoretical basis for the illustration, just the author's limited experience).

For obvious reasons, as well as counting the number of processes/systems for which RTO/RPO have been defined, it is also necessary to check whether the processes/systems actually meet those objectives. A further refinement is to measure the proportions that have been (a) tested or exercised within the past year, regardless of the outcome, and (b) proven, certified and/or accepted by local management as meeting the defined requirements (*i.e.* Actual Recovery Time and Actual Recovery Point less than RTO and RPO respectively). These figures should be higher for tier 1, ideally of course 100%.

Contingency cost metrics

Contingency arrangements can be very expensive so it pays to keep an eye on expenditure. Major cost elements vary from organization to organization but most are likely to include the following:

- IT Disaster Recovery, often contracted to one or more DR specialists and hence incurring an annual charge for a certain level of service;
- Redundant IT equipment, communications links *etc*. (these costs are partially offset by the additional capacity to cover peak loads, but all capacity that is used routinely should be charged to routine expenses and not contingency);

- Contingency planning activities *e.g.* business impact analysis, plain maintenance and exercises. The costs will mostly be man-hours but there are often additional expenses for consultants.

Management confidence metrics

Management confidence in contingency controls can be surveyed through questions such as this:

How confident are you that our contingency arrangements meet the business needs? Please mark the following percentage scale at the appropriate point, in your opinion.

0% 50% 100%

|--+---|

Not at all. Not quite enough | Just about enough Absolutely!

Comments *e.g.* what led you to this score? Have there been particular situations or incidents that influenced your decision?

Here are some supplemental survey questions to tease out more details:

- Do you feel we spend too little, the right amount, or too much on contingency arrangements?
- Relative to our competitors, would you say we are worse, about the same or better at contingency?
- Have we got the right balance between continuity (making sure critical business processes and the IT systems supporting them keep running despite most incidents) and recovery (re-starting processes on standby systems following serious incidents)?
- Which elements of contingency, if any, would you say we need to improve (*e.g.* contingency planning, business impact analysis, contingency exercises, business continuity planning, business process resilience, IT resilience, IT disaster recovery, dual-live/hot-site, warm/cold-site, alternative offices or something else)?

Conclusion

The suggested metrics are intended to help you derive creative and useful measures for your own situation. Do not underestimate the value of presenting and discussing metrics with management. The dialogue can be very effective at teasing out any underlying issues and concerns on both sides, and it promotes adequate investment in information security.

For more information

Please visit Information Security's intranet Security Zone for further information. Additional security awareness materials and advice on this topic are available from the Information Security Manager. NIST's Special Publication 800-55 "Security Metrics Guide for Information Technology Systems" includes an extraordinarily comprehensive list of possible metrics. Andrew Jaquith's book "Security Metrics" is a more pragmatic guide.