Incident Response
From the opening moves to the close

Mel Gosling explores the progression from the initial response to the close of an incident, and finds similarities with chess.

Developing plans to successfully respond to incidents is at the heart of business continuity, but how, when every incident is different, do you plan for the progression of the response as the situation develops over time from the initial response to through to the recovery and return to 'new business as usual', and how and when do you bring the incident to a close?

Because all incidents are different, each response to an incident response is unique. However, there are some things that all incidents have in common, and one of these is a natural timeline. Examining this timeline enables us to identify three phases that we can use to structure our response and cope with the complexity of a developing incident:

- Initial response
- Continuity
- Recovery

These three phases are artificial constructs without well defined start and end points, and the nature and length of both the phases and the timeline varies considerably from one incident to the next. A similar idea is found in the game of chess, which progresses from the opening moves through the middle game to the end game. A plan to win a game of chess consists of an opening gambit followed by using appropriate tactics within a set of pre-defined strategies in the middle and end games, and much the same applies in successful incident response.

Unlike the end of a game of chess, identifying the end of an incident is fraught with difficulty, particularly as some incidents can have long lasting effects on an organisation that do not become apparent for many years. There is therefore, no natural end point to an incident. However, there is a need for an incident to be closed so that response teams can be stood down and a review can be undertaken to identify lessons learned. As such, each organisation needs to set its own criteria for formally identifying the end of an incident, and deciding when it can be closed.

An example of the sort of criteria that might be used would be to declare an incident to be at an end when a new norm of service has been identified and achieved such that there is no longer any need to use specially prepared plans to manage the response to the incident. The logic behind this is that the incident response starts when the specially prepared plans are activated and should therefore end when these plans are no longer needed, but like the three phases of response, this is an artificial construct to enable us to simplify reality.

Like many chess players, business continuity planners tend to concentrate on their opening gambits as these are the easiest to identify and plan for, with the result that most business continuity plans will cover only the first few days or weeks of the response, and not address

the later phases or how to progress from one phase to the next. I refer to this as the organisation's planning horizon.

The majority of incidents that affect an organisation's operational capability are short in duration, such as a power cut, temporary exclusion from an area, or bad weather. As a result, the planning horizon used by most organisations in developing their business continuity plans is perfectly adequate. For longer duration incidents though, such as a major fire, limiting the planning horizon buys time to enable longer term plans to be created and then implemented.

A vital element of any business continuity plan must therefore be to evaluate the incident and estimate the likely duration of its effect on the organisation. If the estimate of the duration is less than the planning horizon then the plans are adequate, if not, then new longer term plans need to be created and implemented before the end of the planning horizon. The shorter the planning horizon the more difficult this is to achieve – which is something that you should consider before determining what your organisation's planning horizon should be.

Deciding on a short planning horizon is perfectly understandable because of the exponential increase in the number of potential situations that an organisation could find itself in as time progresses following an incident. This means that the longer the planning horizon the less useful the plans are likely to be. By following a limited horizon plan the organisation tries to put itself in a favourable position to take advantage of the situation that develops and then go on to a successful conclusion when the incident is closed, just like a chess player moving from the opening gambit to go on to winning the game.

A chess player is a single person who analysis the situation, decides on the most appropriate strategy to deploy, and ensures that the right tactics are used. Organisations though, are made up of many different people and teams. If the incident response is to be successful it needs to be carefully coordinated as if it came from a single mind. This requires a well-managed response at the three fundamental levels at which all organisations operates:

- Strategic – where policy is set and decisions on direction are made
- Tactical – where processes are managed
- Operational – where activities are undertaken

The roles and responsibilities of the recovery teams at each of these levels need to be defined, including when something is outside the authority of a team and needs to be escalated. In this way, any decisions to create and implement new longer term plans when the duration of the incident looks likely to exceed the planning horizon can be taken at the most appropriate level. Minor incidents should be able to be managed at the operational level, but the more significant the incident the more likely it is that the response will be coordinated by the strategic level.

Similarly, if the results of decisions made and any issues encountered are to become known and understood as if there was a single directing mind, then the roles and responsibilities of the recovery teams need to encompass how, where, and when information is escalated from the operational and tactical teams. One of the most common causes of incident response problems is a failure to escalate information in a timely manner to the appropriate team.

The key in all of this is to concentrate on how to minimise the impact of the incident on the organisation, and not to worry about moving from one phase to the next. The progression through the three artificial phases of the response should well defined in procedures and action plans and become seamless when activated. In terms of a preparing your business continuity plans to achieve this aim, you need to:

- Select an appropriate planning horizon
- Develop plans at the strategic, tactical, and operational levels
- Clearly define the responsibilities, authorities, and escalation procedures for each recovery team
- Include the progression as part of the responsibilities, and link it to the planning horizon

Then, when an incident occurs, the strategy to follow from the initial response to through to the recovery and return to 'new business as usual' is:

- Activate the appropriate business continuity plans – the opening gambit
- Estimate the likely duration of the incident – is this less than our planning horizon or are we going to have to plan for a longer duration disruption?
- If necessary, create longer term plans – the middle game
- Either close the incident or implement the longer term plans – the end game

On a final note, although chess is complex, responding to an incident is far more difficult as there are more pieces and types of pieces on the board, the board itself is much larger and multi-dimensional, and unlike chess, the number of things that can happen is not a finite. However, one of the underlying principles of chess strategy is to take control, which is also an underlying principle of incident response strategy. You need to be control of what is happening, moving at the pace of the incident, not at the pace of your decision making.