A Business Continuity Management Simulator

W. J. Caelli, L. F. Kwok, D. Longley

A Business Continuity Management Simulator

W.J.Caelli[1], L.F. Kwok[2], D. Longley[3],

[1] Information Security Institute, Queensland University of Technology, Brisbane
[2] City University of Hong Kong
[3] International Information Security Consultants Pty Ltd.

Abstract. Comprehensive BCM plan testing for complex information systems is difficult and expensive, if not infeasible. This paper suggests that a simulator could be employed to ameliorate these problems. A general model for such a BCM simulator is presented, and the implementation of a prototype simulator is described. The simulator reacts to system disturbances by seeking alternative configurations provided within the BCM plan, reporting the resource availabilities in the updated system and identifying any failure to meet the requirements placed on the system. The simulator then explores any changes in data security introduced by the proposed post disturbance configuration and reports any enhanced risk.

Keywords: BCM planning, simulator.

1 Business Continuity Planning

In a survey of 94 Australian organizations in 1999-2000 [1, 2] the majority of organizations stated that the longest time they could be out of action was less than 24 hours. Moreover 30% of these organizations said that their longest out-of–service time was less than 8 hours. It is perhaps sobering to reflect that in the subsequent 9 years the world has experienced major terrorist attacks in capital cities, invasion of Iraq, two tsunamis, the worst Chinese earthquake in a decade and devastating weather events worldwide all having implications for BCM. Given society's increasing dependency on distributed, complex information systems it might appear the risks associated with loss of availability in such systems have both enhanced consequences and likelihood.

Traditionally organizations developed Disaster Recovery Plans to ensure that alternative facilities were available if some major event impacted upon their mainframe computer systems. The development of complex, distributed systems combined with organizational reliance upon on-line operations emphasized the importance of business continuity management which seeks to minimize the likelihood and magnitude of potential business interruptions, and encompasses Disaster Recovery Plans to guard against the major loss of IT services at any level in a system hierarchy.

The Australian National Audit Office guide on Business Continuity Management states the objective of business continuity management is to *ensure uninterrupted availability of all key business resources required to support essential* or *critical*

business activities. One may query if this is a sufficiently comprehensive definition of the objective since there is no mention of continued compliance to required governance and security policy. Presumably a bank would not be satisfied with a business continuity plan that left a security loophole for major fraud in its fall back procedures.

Australian Standards HB292-200 [3] provides comprehensive guidance to practitioners on the establishment of effective business continuity management within organizations. The guide emphasizes the close relationship between risk management and business continuity planning. In particular risk management assists in the identification of cost effective controls to minimize the likelihood of a business interruption, whilst business continuity plans will, inter alia, addresses the actions required to deal with significant disruptive events. The guide does not however appear to address the problem of risk analysis for the systems in the post disruption phase and again there is no mention of the requirement for continuous conformance to security / governance policy.

This paper thus addresses two issues of BCM planning: availability testing, and identifying enhanced risks following a disturbance. Simulators are commonly employed for training on complex systems in situations where real experimentation would be hazardous or infeasible, and would therefore appear to be of value in testing defenses against critical, complex information system downtimes.

A proposed model applicable to BCM planning in a variety of systems is introduced, and the implementation of such a simulator developed for information systems is described. The simulator reacts to system disturbances by seeking, where necessary, alternative configurations provided within the BCM plan, reporting the resource availabilities in the updated system and identifying any failure to meet the requirements placed on the system. The simulator then explores any changes in data security introduced by the proposed post disturbance configuration and reports any enhanced risk.

2   BCM Simulator Design

2.1   Overview

The BCM problem is illustrated in Fig 1. The system under review receives resources from external sources and in turn provides a set of resources for the external world, a subset of these resources are deemed essential in all circumstances. An interruption in the supply of external resources, and/ or some internal disturbance to the system, will have an impact upon the system's ability to supply the requisite resources.

BCM planning will provide some inbuilt defences against such external disturbances; however the system is dynamic and will therefore exhibit a transient response to the disturbances. The BCM Planner may be called upon to report upon the response to specified disturbances in terms of:

- time duration of transient response;
- system output during the transient response;

- system output in the subsequent steady state;
- effectiveness of security systems in the transient and steady state.

At this stage the simulator deals only with steady states, i.e. it does not deal with any potential delays in providing alternative sources or in any race conditions that may arise in (say) invoking load sharing algorithms.

When a disturbance causes a source to fail, and one or more alternative sources are provided by the BCM plan, the simulator incorporates these backup sources in preference order. Any shortfall in these alternative sources will have consequential effects on the system's ability to supply essential resources and these consequential effects are also reported. The simulator then explores potential changes in the data security scenario of the post disturbance configuration and reports any risk changes.

2.2 Model

Overview. A top down view of BCM suggests that the system under review is in effect a resource transformer subject to disturbances. The system imports resources from the outside world and converts them into exported resources. A subset of these exported resources is deemed essential and must be maintained in spite of interruptions to imported resources or disturbances to the system. This is the conventional view of BCM but with current complex information systems we should also add the requirement that the system should maintain the specified system security policy under all circumstances. Drilling down into the system will reveal some network of resource interconnected subsystems. The original set of imported and exported resources will be enhanced by internal resources generated by the subsystems. Looking in more detail at the operation of the subsystems we may further specify the subsystem as a collection of sink/source units. Each such sink/source unit provides a source for a particular resource type, e.g. bank transactions; a collection of sinks associated with the source specifies the imported resources, in type and quantity, required to guarantee the supply from the associated source. In a BCM model the nature of the subsystem transformer actions is not relevant to the study. The essential point is that the sources require certain satisfied sinks to ensure their continued operation. Some subsystems may be considered from the system viewpoint to be pure sources or sinks, e.g. the system imported resources are derived from pure sources and the exported resources are collected by pure sinks (See Fig 1).

The BCM plan is concerned with the operation of the system following an interruption to an imported resource or a disturbance. In the terms of the model as described above external resource interruptions are modelled as a disabling of external sources, and disturbances similarly disable internal sources. From a BCM viewpoint there are two significant facets of this situation:

- the alternative supplies to internal sinks, i.e. backup sources, to ensure that external sinks deemed essential continue to be supplied;
- the relationship between a disturbance and the failed internal source.

A system with no source redundancy would fail to export one or more resources in the case of any external source interruption or system disturbance. BCM plans

therefore incorporate source redundancies to maintain the export of essential resources. This redundancy implies that some sinks are connected to multiple sources with various preference levels (See dotted link in Fig 1).

At this stage the model relates to any definable system and the simulator was developed to be context independent as far as possible. However in order to move to a prototype stage it was decided to concentrate initially on information systems so that ideas could be tested. Having described the system in terms of subsystems the next stage was to represent subsystems as interconnected components, or *entities*. In information systems these entities comprise:

- computers: servers, gateways, workstation clusters;
- locations: sites, buildings, floors, rooms;
- services: power supplies, cabling;
- switches: for physical resources, e.g. power, human resources;
- data networks.
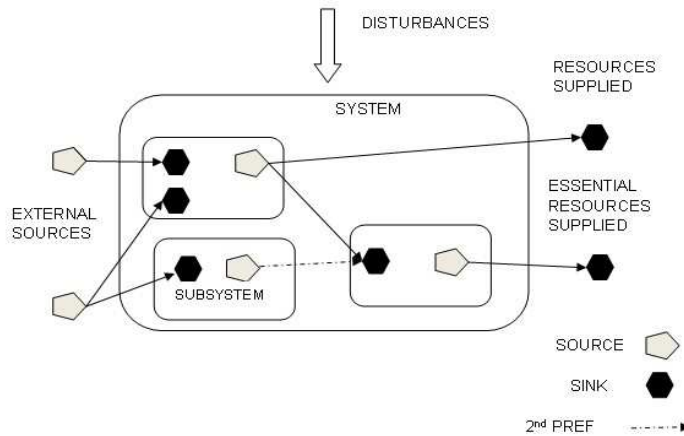
Fig. 1. BCM Model Subsystems are interconnected by resources and comprise sinks and sources. Some sinks may be supplied by two sources with different preference levels.

With this component definition the relationship between a disturbance and the failed internal source can now be addressed. Fig 1 illustrates subsystem with apparently isolated sinks and sources. The simulator *entities* have a more structured set of sinks and sources:

- an entity can have multiple sources, each for a unique resource, e.g. a server can be a source of multiple data items;
- a source produces the result of a transformation of other imported resources, the nature of the transformation is irrelevant in BCM terms and can simply be formulated as a set of sinks feeding each source.
- the entity is a physical reality and the functioning of that entity is dependent on certain physical resources, e.g. a server's continued operation depends upon the supply of electrical power, manual operators, maintenance technicians etc. Any

significant interruption to these resources will result in the disabling of all other sources in that entity, i.e. a server will not export files if it has lost power.

If an entity does not receive the set of resources necessary for its physical operation, then all the logical sources in that entity will fail. This situation is represented in the simulator model with an artificial source, termed an *entitysource* which supplies each entity source with an artificial resource termed an *entityresource*. The *entitysource* has a sink for each imported resource required for the physical operation of the entity, and each entity source has a sink for such an *entityresource*.

Hence the relationship between a disturbance and internal sources can be formulated:

- a disturbance inhibits the flow of resources to *entitysource* sinks;
- the *entitysource* with at least one unsatisfied sink is disabled;
- the remaining entity sources now have their *entityresource* sink unsatisfied and are themselves disabled (See Fig 2).
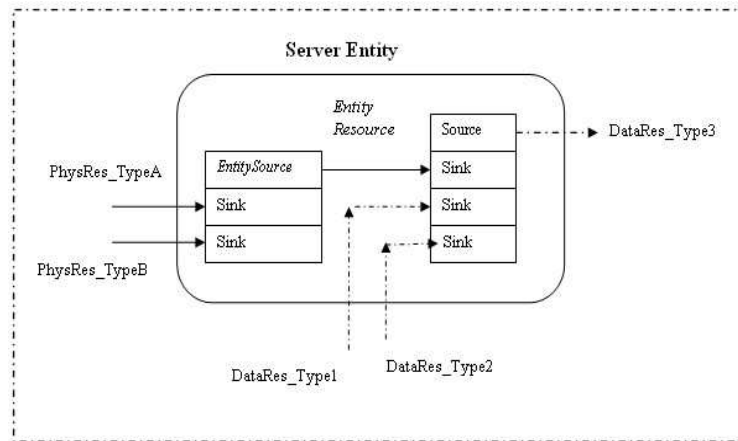
Fig. 2. Sources and Sinks in an Entity. If *entitySource* fails, due to loss of any physical resource, then all other sources within the entity are deprived of the *entityresource*, and fail.

2.3   Simulator Resources, Entities and Interconnections

Resources. The simulator is designed to demonstrate the flow and interaction of resources through the system after various types of disturbances. The nature of these resources significantly influences the design of the simulator model and in the context of an information system the resources were categorised as physical or data. Each resource is described in terms of type and quantity, but there are major differences between the roles of quantity in each category.

Physical resources are heavily quantity dependent, e.g. a power resource may be described as electrical power 240V AC (type), and a numerical value for wattage

(qty). Similarly a server operator resource may be specified as Operator Grade_1 (type) and a number of full time hours (qty). If the quantity demand for a resource exceeds the maximum available quantity supply then the supplying source will be overloaded and some sinks become unsatisfied. If a physical resource sink is unsatisfied then the simulator seeks an alternative source, if no such alternative is found the consequent effects of the unsatisfied sink are explored, e.g. sources in the same entity are disabled producing consequential unsatisfied sinks elsewhere.

Data resources are similarly specified via both type and quantity, but data sources are not so susceptible to overload. An excessive demand for bytes from a server, or excessive traffic on a network will involve performance issues, and are reported by the simulator, but they will not lead to a source overload and disablement, with consequential unsatisfied sinks. The originating source also plays different roles in the two categories of resources. A physical source takes the view that any power supply is a compatible supply. However a data resource sink will specify both the resource type and the source for a data resource, i.e. a user will connect to a specific server in order to download a file with a given name.

An information system will commonly have a vast multitude of data resources but from a BCM viewpoint they may be aggregated. For example if a server normally collects data from one other server, but has backup servers each supporting part of that total data set, then the data resources may simply comprise the aggregated data sets from each of the backup servers.

Entities. All entities contain an entitysource (See Fig 2) and one or more other source/sink units, but the nature of these sources and sinks varies with their category.

Computers. Computers in their various roles have an *entitysource* with physical resource sinks, e.g. electrical power, plus one or more source/sink unit importing data resources and exporting some transformation of the input data.

Location. From a general BCM viewpoint sites and buildings serve as potential sources of access control, protection for equipment against severe weather events etc. From the BCM simulator model viewpoint the hierarchy from sites to rooms, cupboards etc., also provides a potential conduit for physical and data resources. A server located in a room may be deemed to be served by the resources available on that site: power, trained staff etc. This conduit concept reduces the impact of a major problem for BCM simulator users, i.e. specifying the multifarious interconnections between sources and sinks. For many resources, location entities may be considered as special types of switches (See Switches below).Once a server is allocated to a room (See Relationships) the BCM simulator software can automatically create the links between the server sinks and the room switch sources.

Switches. A cabling system is a typical switch, receiving power from some preferred supply source and delivering it to various computers etc. The source /sink unit differs from that of a computer source/sink inasmuch as the source resource has the same type and quantity as that of its single sink. If the switch has a number of alternative supply sources this sink will be connected to each in a specified preference order. The source also has an attribute indicating the maximum safe loading.

Data Networks. Data networks serve as switches for data resources and differ substantially from the switches described above (See Fig 3). In the BCM model a data network provides proxy sink/source units for each data resource transmitted through the data network. These proxy units also contain the name of the entity that hosted the original source for the data. As discussed in Locations above the BCM simulator software develops the source – sink links from originating source to destination sink.

Networks are similar to switches inasmuch as an excessive volume of traffic may affect performance, since there are multiple data sources the congestion level is an attribute of the network *entityresource*.

Fig. 3. Network serving Server_1 and Server_2 and source-sink interconnections. The network includes proxy sources corresponding to those of both servers.

Interconnections. The source-sink links are so numerous, for even simple systems, that the software was designed to minimize the manual effort of link specification.

A source sink interconnection in the BCM model is normally specified by a relationship linking the host entities, internal entity source – sink links are automatically inserted by the software. The relationship specified by the model builder contains sufficient information for the software to identify the particular source and sink in the host entities, and the preference level of the link.

Relationships for entity – data network connections, can be supplied with details of the desired data source, and preference levels if more than one such source is specified. Similarly if a data network is to be connected to one or many other networks, the preference levels are specified.

2.4 Simulator Operation

Overview. The BCM simulator is designed to provide the user with various system scenarios following specified single or simultaneous multiple disturbances. The simulator seeks to mitigate the effects of the disturbances by seeking alternatives for any sources disabled by the disturbances, and then displays the post disturbance scenarios. The simulator also provides information on any additional risks in this post disturbance scenario.

The simulator operates in two phases: setup and interaction. The setup phase itself has three stages. First the simulator employs all the user specified data and constructs the model, much of this construction activity is devoted to establishing the multifarious source-sink connections throughout the system. The second stage then checks for any sources providing no resources, and for any subsequent unsatisfied sinks. When the most preferred source for a sink is unable to supply the required sink resource quantity, alternative sources are sought according to specified preferences. If these attempts fail, then the unsatisfied sinks cause the associated sources to be disabled and the simulator explores and attempts to rectify these consequential resource shortfalls. Finally, in the third stage, the simulator checks and reports the risk levels of all entities.

In the interaction phase the user can explore the effect of specified disturbances and observe graphs of resource flows through the system. Enhanced risks for entities are colour coded in these graphs.

Responding to Loss of Availability. If the total load required by sinks from a source exceeds the quantity available at that source then alternative sources are sought according to the sink preference levels. The source overload could be addressed with load sharing, i.e. switching some sinks to alternative sources until the source load is reduced to its available capacity. The simulator currently employs no such load sharing algorithms. If a source is overloaded it reacts like a cut-out in an electrical system and is disabled. The full load demanded from that source is then directed to the specified alternative sources. The simulator in effect warns the user of source overloads, but any load sharing must be performed manually, i.e. the user recommences system setup with changes in some sink preferences effectively diverting the load to other sources.

Risk reporting. In highly complex information systems employing inbuilt safeguards against various loss/denial of services, there are two post disturbance considerations:
- is the set of safeguards sufficient to guarantee essential services following one or more particular disturbances?
- is the system security policy maintained in the new system configuration?

The previous sections have dealt with the first issue. In this section any enhanced risk associated the post disturbance system configuration is discussed. Risk reporting differs from problems associated with availability levels in that there is greater uncertainty associated with the consequential effects of the new risk. If a building suffers physical damage a server located in a room of that building may or may not continue to function. The simulator thus restricts itself to reporting the potential

spread of such risk but it does not invoke a search for alternative less risky configurations.

Data security is complex and may well be impacted by configuration changes following a disturbance. For example, a VPN network fails and sensitive traffic is switched to a network vulnerable to data eavesdropping. Encryption itself can be a source of risk if alternative sources are switched in, e.g. if a sink is setup to check data integrity it will reject data from an alternative source that provides no such integrity field.

The BCM simulator checks data security risk over the data path from source, through any intervening networks, to sink and reports on the risk associated with the path, e.g.

- data passing though a congested network;
- unencrypted data in a network with a disclosure hazard;
- any encryption incompatibilities encountered in the path.

A number of physical risks may also be reported, for example if the total quantity supplied from a switch source exceeds the switch safety rating, then this risk is reported. The risk reported is not a probability – consequence measure but simply indicates a situation has arisen that will be of interest to the user.

Fig. 4. BCM Simulator Graph. Path of Profit_Loss_File from Acct_File_Server to Exec_File_Server. Org_WAN is down and traffic is diverted to Pub_WAN that has a disclosure risk reflected in risk at two servers.

The risk is currently coded as low, moderate or extreme, An extreme risk is associated with loss of availability of an entity, risk relating to a resource, e.g. data resource transiting a congested network, is classed as moderate. The risk reporting involves colour coding the appropriate entity on the graph (See Fig 4).

3   BCM Simulator Implementation

The simulator was implemented as an extension of the ISM (Information Security Model) software described in previous papers [4, 5]. As such it demonstrated the value of security documentation developed in a database form and used to support various security management tasks, such as risk assessment, compliance audits, BCM, etc.

The model entities correspond to the ISM entities and the BCM data was simply added to the entity data. The graph facilities for illustrating threat networks were adapted to the demonstration of resource paths (See Fig 4). The simulator software as described above was then included in the ISM package.

4   Conclusion

This paper describes the experience gained in the development of a prototype business continuity management simulator. The simulator software was an extension of a package originally developed for risk studies and automated compliance testing. The business continuity management simulator software relied heavily on the facilities provided by the earlier package to search for system entities in a tree hierarchy to represent ad hoc relationships between such entities, to specify complex attributes and to provide graphical representations of system behavior. This justified earlier confidence in the versatility of the ISM package.

This study has demonstrated the feasibility of developing a simulator which can bridge the gap between paper/team based and live scenario exercises in business continuity management training and planning.

References

1. Musson, D. and Jordan, E.: *Managing for Failure – The Macquarie University Survey of Business and Computer Contingency Planning in Australia*, Macquarie Graduate School of Management, Australia. (2002).
2. Australian National Audit Office: *Business Continuity Management*, August (2000).
3. Australian Standards HB292-200: *A Practitioner's Guide to Business Continuity Management*, January (2006).
4. Kwok, LF. and Longley, D.: Security Modeling for Risk Analysis, *Proc. 18th IFIP World Computer Congress*, IFIP 2004, 22-27 August 2004, Toulouse, France, pp29-45, (2004)
5. Branagan M., Caelli W.J., Lam-for Kwok, Longley D.: Feasibility of Automated Information Security Compliance Auditing, *Proc. IFIP TC 11 23rd Int. Information Security Conf.*, September 2008, Milan Italy, pp493-507, (2008).