Microsoft®

MOF Service Management Function Service Continuity Management

patterns & practices

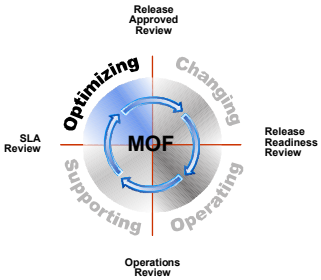

Microsoft®
Solutions for Management

Contents

Document Purpose

This guide provides detailed information about the service continuity management service management function (SMF) for organizations that have deployed, or are considering deploying, Microsoft technologies in a data center or other type of enterprise computing environment. This is one of the more than 20 SMFs defined and described in Microsoft® Operations Framework (MOF). The guide assumes that the reader is familiar with the intent, background, and fundamental concepts of MOF as well as the Microsoft technologies discussed.

An overview of MOF and its companion, Microsoft Solutions Framework (MSF), is available in the *Introduction to Service Management Functions* guide. This overview guide also provides abstracts of each of the service management functions defined within MOF. Detailed information about the concepts and principles of each of the frameworks is also available in technical papers available at www.microsoft.com/solutions/msm/.

Executive Summary

In today's highly competitive and service-oriented business environment, companies are judged on their ability to continue to operate and provide a service at all times. This availability is accomplished through a balance of risk-reduction measures such as resilient systems and recovery options including backup facilities. Successful implementation of contingency management can be achieved only with visible senior management commitment and the support of all members in the organization.

Ongoing maintenance of the recovery capability is instrumental if the solution is to remain effective. This is achieved through:

- A rigorous configuration, change management, metrics monitoring, availability management, capacity management, and review process.
- Education, documentation, and awareness for the whole organization and its customers.
- Specific, ongoing training for personnel involved in the process.
- Regular testing and analysis of contingency and disaster recovery plans.
- Vigilance to change, and the introduction of new risk.

Many administrators believe if their data centers are not in a flood or earthquake zone, the need for a catastrophe solution is slight. In reality, this is not the case. A natural catastrophe can occur in a not-so-natural way. For example, power grids routinely experience random brownouts and blackouts. The loss of power for a corporate enterprise site can mean a great number of lost productivity hours and revenue.

Service continuity depends on a successfully thought-out contingency plan that is tested and evaluated on a periodic basis. A contingency plan, no matter how well thought-out and tested, does not guarantee service continuity. It is important to incorporate a structured daily approach to maintaining a corporate enterprise infrastructure and to ensure that proper procedures and policies are put in place and followed. Once these procedures and policies are put in place, vigilance to change and periodic re-evaluation of the proposed contingency plan is essential to maintaining the validity of the plan.

Process and Activities

Service Continuity Management Overview

The objective of the service continuity management process is to ensure that any given IT service is capable of providing value to the customer in the event that normal availability solutions fail.

Availability management and service continuity management have become two of the most importance aspects of service delivery in the highly visible e-commerce global economy. The demand for operations to be available 24 hours a day, seven days a week (24x7) is greater than ever. Availability, or the lack of it, has a dramatic influence on customer satisfaction and can quickly impact the overall reputation and success of the enterprise.

Many factors affect the availability of an IT service such as hardware failure, environmental issues, and human error. A hardware failure, such as a broken power supply or disk drive, is one of the most obvious factors to consider. If the only power supply to a server fails, then this might cause the whole IT service to be lost. Dual redundant power supplies attached to the server can be employed to mitigate some of this risk. If power to the whole computer room or data center is disrupted, battery backup systems can be employed to cover the short time it might take to start up a standby generator. Exposures such as these are referred to as availability risks and the actions we might take to mitigate them are called countermeasures.

If any of the countermeasures fails, then more drastic actions must be taken. These actions are outlined in a document called the contingency plan.

For example, one data center has a standby generator that is designed to take over if utility power is lost for any reason. During a particularly cold winter day, ice snapped the power lines to the building. The operations manager tried to start the generator but the diesel fuel had turned to slush from the cold. In this case, an availability solution existed, but it failed. Therefore, the manager had to implement the contingency plan instead and move the service to another site.

Risks to availability also exist within processes and procedures and also arise out of human error. If a poorly tested change is introduced that inadvertently prevents users from connecting to the IT service, then the complete service is unavailable until access is restored. If the production database is accidentally overwritten with last night's backup data instead of a new backup being performed, this can have catastrophic consequences on availability. Countermeasures, such as carefully designed testing and release procedures, and appropriate staff training plans, can also be employed to help mitigate these risks. When these mitigation plans fail, contingency plans must be invoked.

Availability management and service continuity management are closely related in this respect as both processes strive to eliminate risks to the availability of IT services. The prime focus of availability management is handling the routine risks to availability that can be reasonably expected to occur on a day-to-day basis. Where no straight-forward countermeasures are available or where the countermeasure is prohibitively expensive or beyond the scope of a single IT service to justify in its own right, these availability risks are passed on to service continuity management to handle.

Goals and Objectives

Service continuity management is concerned with managing risks to ensure that an organization's IT infrastructure can continue providing services in the event of an unlikely or unanticipated event. This is accomplished through a process that analyzes business processes, their impact on the organization, and the IT infrastructure vulnerabilities that these processes face from a myriad of possible risks. This requires a great deal of research to be conducted with diligence to identify all critical business processes and their vulnerabilities.

This task begins by dividing the effort into three phases: define the service level objectives, propose a solution to meet those objectives, and formalize the written agreements and contingency plan. Each phase has tasks and deliverables associated with them that assist in determining cost-effective solutions. The deliverables need to be maintained as active documents and updated as needed.

Scope

Service continuity management primarily considers those IT assets that support key business processes. However, the installation of mechanisms to deliver service continuity management will not necessarily be sufficient to keep those business processes operating after a service disruption. Should it be necessary to relocate to an alternative working location, provisions are required for items such as office and personnel accommodations, copies of critical paper records, courier services, and telephone facilities to communicate with customers and third parties.

The service continuity management process identifies the required and agreed minimum level of business operation following a service disruption, along with a requirements definition covering systems, facilities, and service requirements. The process then examines the risks and threats to these requirements and develops an IT risk reduction or mitigation program. This program implements mechanisms delivering the continuity requirements necessary to provide the required optimum level of business operation.

Key Definitions

Cold site, fixed center. This can include the provision of empty accommodations fully equipped with power, environmental controls, and local network cabling infrastructure, telecommunications connections, and available in a disaster situation for an organization to install its own computer equipment.

Cold site, mobile center. This option is the same as a cold site, fixed center; except the site is mobile or portable. This site may be erected on a pre-designated location, or near the actual facilities.

Warm site, fixed center. A location with suitable computer equipment ready to recover service.

Warm site, mobile center. Commercial recovery services can be provided in portable form where a pre-configured computer is delivered to a customer's site, within a certain time; typically 24 hours. The computer equipment is contained in a trailer and transported to the site by truck. The trailer is outfitted as a computer environment with the necessary services and only needs power and telecommunications links from the site to the trailer for the service to be established.

Hot site, fixed center. Dedicated computer equipment mirroring critical business systems ready to take over immediately with no loss of data.

Business impact analysis. A business impact analysis (BIA) focuses on the business needs of IT services. Being without any IT service will have a detrimental effect on the business, but the severity of the impact will vary with time and also be affected by its point in the processing cycle. The impact in the loss of a real-time service, such as trading in a money market, will be felt immediately while the business may cope for some time without other services. While establishing the urgency of each service, the BIA identifies the minimum requirements of each service to meet the critical business needs.

Contingency plan. A tested plan, documenting the actions to be taken and implemented in the event of a disaster.

Major Processes

Service continuity management comprises of four main processes and a number of sub-processes as follows:

- Acquire service level requirements
    - Identify information technology service layers
        - Service
        - Application
        - Middleware
        - Operating system
        - Hardware
        - Local area network
            - Passive components
            - Hubs, switches, and routers
            - Network interface cards
        - Facilities
            - Edifice
            - Environmental controls
            - Physical security
            - Fire suppression
            - Human convenience

        - Egress
            - Security
            - Water
            - Sewage
            - Gas
            - Electricity
            - Internet access
        - Identify risks to each information technology service layer
    - Propose contingent solution
        - Design for failover
        - Outsourced services
        - Facilities
            - Cold site, fixed center
            - Cold site, mobile center
            - Warm site, fixed center
            - Warm site, mobile center
            - Hot site, fixed center
        - Personnel
    - Design for recovery
    - Designing for customer satisfaction during outages
- Formalize operating level agreements
- Formalize the contingency plan
    - Definition of contingency levels
    - Escalation and notification procedures
    - Startup and shutdown procedures
    - Communications methods
    - Status reporting requirements

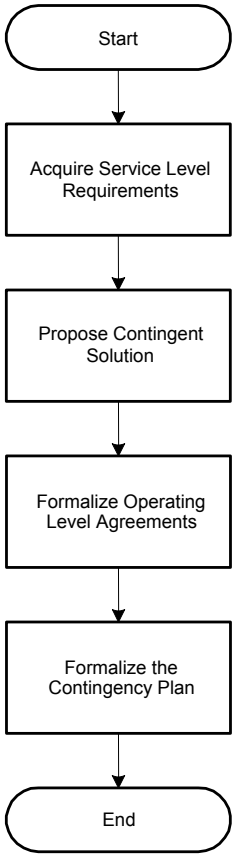

Figure 1
Service continuity management process flow

Acquire Service Level Requirements

The availability management SMF defines availability requirements for business services. Service continuity management typically continues where availability management left off to address those availability risks that the availability management SMF cannot or chooses not to address.

Once risks are known, the users, with the help of IT, must decide which risks are to be mitigated and which are to be assumed. Mitigating a risk requires people, time, and money. IT management might determine that a risk is so small that they do not want to incur the cost of mitigating it. For example, meteor damage to the data center often falls under this category. The likelihood of a meteor shower damaging the data center is so small and the cost of installing anti-meteor shielding to the building is so large that most organizations choose to assume the risk rather than assuming the cost of mitigating the risk.

The risk of a meteor shower destroying your data center is a possibility. Just as the risk of tornadoes, floods, and any number of unknown events are also possibilities. It is probably not cost effective to create custom solutions for every eventuality. It is much easier to create a single plan that can be implemented in any of these events. This plan is called the contingency plan.

Service continuity management starts by carefully agreeing to availability targets with the customer and determining the cost of downtime or unavailability of the IT service in question so that a realistic IT budget can be established. It is also important that the negotiation include realistic expectations of reduced system availability while the contingency plan is in place. This process involves an element of education and negotiation on both sides (the customer and the IT organization). The customer needs to understand how to define and articulate their availability requirements. The IT organization needs to understand the functions that make up the overall IT service and which of them are the most critical.

Identify Information Technology Service Layers

To understand where risk may be introduced, the IT environment must be broken down into logical, manageable components. One way to do this is to divide the services provided by the IT department into layers.

For the IT department, these layers are:

- Service
- Application
- Middleware
- Operating system
- Hardware
- Local area network
- Facilities
- Egress

The operating system requires hardware on which to run. In turn, the operating system provides critical functions to middleware services, such as a database. Information technology cannot deliver a service (top layer) unless all of the services underneath are functioning properly. The key to providing a service, then, is to provide the supporting IT functions required to create that service.

Service

This is the function that IT is helping the business perform. This likely has a name that is easy for the business to understand, such as payroll. In order to perform this function, the business needs the support of the following IT layers:

Application

The application is the top-most layer of the IT stack. The application is the piece of IT that most users see. In order to provide a service, such as accounting, the users make use of an application, such as Microsoft® Great Plains® Small Business Manager.

Some IT organizations might consider the Web services, such as Microsoft® Internet Information Server (IIS), as the application and would not consider static Web site content an application since it does not "run." To some extent the distinction is arbitrary; people see the Web pages but most are not aware of the services required to deliver them.

Middleware

Middleware can be defined as "the part of the application that the users do not see." Middleware includes databases, Web services, and messaging systems. Given this broad definition, the question of what constitutes middleware varies greatly from application to application. In every case, however, a complete map of the middleware should be made so that targets for its availability and capacity can be accurately created.

Operating System

The operating system is the software that controls the allocation and usage of hardware resources such as memory, CPU, disk space, and peripheral devices. Because of this relationship between software and hardware, correct operating system performance is critical to correct application performance. The user may not be aware of the operating system or what it does, but the user is aware when critical services are not available due to poorly written device drivers and other operating system problems.

Hardware

Hardware, as used here, represents a wide range of component types. While this not a comprehensive list, it certainly contains computers and storage devices. Within each of these there can be memory, fans, power supplies, and many other device types. Regardless of how it is defined, extra hardware and spare parts need to be allocated to a data center in order to replace pieces that might fail.

Local Area Network

The network inside a building provides a communications backbone upon which computer systems are able to communicate. For a more detailed explanation of networking components, see the network administration SMF guide. The network may contain several components including the following:

Passive Components

Passive components, such as wires, wall jacks, and so on, are key to any network. These components must be accounted for in availability calculations since they may occasionally break and therefore need replacing.

With the increasing use of wireless networking technologies; however, the use copper cables is becoming less common.

Hubs, Switches, and Routers

Hubs, switches, and routers are critical to any network infrastructure. They control and route data on the network. Each must remain available, and must provide sufficient capacity to meet the needs of the layers higher in the stack.

Network Interface Cards

Computer systems connect to the network backbone with network interface cards. These components are relatively inexpensive, so it is not uncommon, for example, to place more than one card into each computer to provide redundancy in the event of one failing. It is also common to direct all external Web traffic to one card and reserve the other for internal management data.

Facilities

Facilities consist of the building that houses the data center and any associated components. These components may consist of:

Edifice

The physical building is obviously very important as it provides a shelter from the elements. A good building provides security for its occupants and a means by which an artificial environment can be maintained.

Environmental Controls

Heat cycling is one of the main causes of failure for any mechanical system. Constant expansion and contraction of metal parts caused by variations in temperature can cause metal fatigue in cases and racks. It has also been know to force electrical components out of their sockets. Therefore, most data centers maintain a constant temperature in order to reduce the stress on the systems caused by heat cycling.

Also of concern is the amount of heat human bodies and computer systems generate. In all but the coldest environments, the heat generated by the computers is sufficient to overcome most outside temperatures. Therefore good heating, ventilation, and air conditioning (HVAC) systems are of critical concern to IT managers.

Physical Security

The best data firewall system is useless if someone can simply carry your systems out through an unlocked back door. The safety of the occupants in the buildings must also be ensured regardless where the data center is located.

Fire Suppression

Fire can devastate a data center in minutes. However, water that is commonly used to fight fire can cause as much damage on the computer hardware as the fire itself. Therefore, systems that remove the oxygen from the air without damaging the computers, such as Halon, are in common use. Depriving the air of oxygen is damaging to humans; however, so adequate escape doors are also required.

Human Convenience

People require facilities just as computer systems do. These services may be as basic as water and toiletries, or they may be as sophisticated as cafeteria and lounge facilities for off-duty personnel. The needs of the human portion of IT should not be overlooked when negotiating agreements with customers.

Egress

Egress can be defined as anything that leaves the facility, or it is connected externally. Primary egress facilities are typically outside of the direct control of the IT department and, therefore, must be provided by another party. Because they are critical, IT may be required to provide secondary sources for all egress services in the event of a failure of the primary system. Egress services typically consist of:

Security

Security to the facility provides a level of access to personnel. This level of access may in some instances introduce risk. Providing for a secure data center, where access is recorded and violation of access also triggers alerts may assist in preventing unauthorized entry, and introduction of risk.

Water

Water is more critical to the people that work in a facility than to the technology components. However, sometimes water is used in cooling towers in order to maintain the environmental controls in a facility.

The availability of clean water is a critical consideration when dealing with site recovery scenarios. It is not enough to simply have new computers delivered to a storm-wrecked data center if the personnel required to install the systems do not have clean water.

Sewage

Like water, sewage removal is critical to the IT staff. This consideration is important when dealing with an emergency or site recovery situation.

Janitorial Services

All companies have waste. Often this waste is of little or no use to anyone. At other times, this waste may be useful for the data it contains. Secure janitorial services ensure that proprietary or dangerous wastes are disposed of in accordance with local environmental laws and without other undue risks to the company.

Gas

In some areas gas is used to provide heat to a facility during cold months. Gas is highly flammable; however, and care must be taken to ensure that a spark from one of the numerous sources of electricity does not ignite a gas leak in a facility.

Electricity

Computers are voracious consumers of electricity. A single data center can consume as much electricity as hundreds of homes. Reliable and plentiful electricity is critical to the correct operation of the technical components of a data center.

Internet Access

Most data centers do not operate in a data vacuum. Data must flow into and out of the center. Therefore, some sort of connection to users or other computing centers is required. This connection may be a privately owned, point-to-point wide area network (WAN) link, or it may be a link to the Internet. In either case, it is likely that an external supplier will provide the network services. For more information on egress Internet access, see the network administration SMF guide.

Identify Risks to each Information Technology Service Layer

By examining risks and vulnerabilities to each level of the IT stack, single points of failure are identified, as well as a single layer that may address risk on the layers above.

At a minimum, the following risk assessment activities need to be performed:

- Identify risks to particular IT service components (assets) that support the delivery process, which causes an interruption to agreed service on each layer in the following scenarios:
  - Damage or denial of access to corporate assets by customers, employees, business partners or operators.
  - Loss of IT systems, networks, private branch exchange (PBX), automatic call distribution systems, firewalls, cryptographic systems, public key infrastructure (PKI), and so on.
  - Loss of customer or internal data or loss of integrity to data.
  - Loss of network services including telecommunications providers.
  - Unavailability of key staff (for example, only one person knowing how to maintain a particular critical network server or business application) and no existing documentation.
  - Failure of partner or service providers (such as support, development, or maintenance).
  - Breach of security (such as fraud, sabotage, computer viruses, or malicious software).
  - Loss of environment (such as air conditioning).
  - Loss of critical paper records or media (such as manuals, documents, backups, and so on).
  - Loss of utilities (such as power, gas, or water).

- Assess threat and vulnerability levels. Threat is defined as how likely it is that an incident will occur and vulnerability is defined as whether, and to what extent, the organization is affected by the threat materializing. A threat is dependent on such factors as:
    - Assess the likely motivation, capability and resources for deliberate incidents such as malicious damage to organization computer systems, commercial failure of a key technology provider, attack against an Internet service providers (ISPs) and active service providers (ASPs) servers, and corruption of ISP/ASP solutions and or data.
    - Assess your service provider's location, environment, and quality of internal systems and procedures (for accidental incidents).
    - Assess your organization's location, environment, and quality of internal systems and procedures (for accidental incidents).
    - Look for single points of failure for the delivery of services. (For example, a travel agent relies on information feeds for flight bookings delivered by an ASP. If the link fails, flights can not be sold.)
    - Assess the levels of risk. The overall risk can then be measured. This may be done as a measurement if quantitative data has been collected, or qualitative using a subjective assessment of, (for example) low, medium, or high. An example of a tabular format used to express the level of risk is shown below. Each risk can be assessed in terms of the associated threat and vulnerability. Using the table, it is possible to determine the probability of specified risks occurring (for example, a high threat and high vulnerability implies a high probability of occurrence).

Mapping these risks against the IT layers above provides a clear picture of where vulnerabilities exist, as well as the impact they will have.

When examining the layers thought vulnerable, solutions can be found on the IT layers below, where risk or vulnerability may have not even been considered. The risk assessment process can also identify risks that are specific to an individual technology, or application. By identifying all dependencies to these processes, the ability to conduct a successful recovery is greatly increased.

The following is a portion of a risk assessment. It is important to understand that this does not address every risk, only a few. Also, the block in which the vulnerability is present renders each block to the right as vulnerable as well. If a facility experiences a fire, the network, hardware, operating system, and so on become unavailable.

Table 1. Risk Assessment

Risks	Egress	Facilities	Network	Hardware	Operating System	Middleware	Application	Service
Fire		Med						
Flood		Low						
Virus					High			
Power outage	Med							
Logon failure			Med					
Lack of staff			Med					
Human error		Med						

All risks to the availability of each IT component need to be considered. The nature of the risks faced by an IT component varies according to the MOF IT layer the component resides in.

Example availability risks by IT layer:

- Application, middleware and operating system layers:
  - Single point of failure
  - Incorrect configuration option
  - Design flaw
  - Poor development methodology
  - Coding error
- Hardware and network layers
  - Single point of failure
  - Out of date firmware
  - Poor documentation
  - Vendor support quality
  - Lack of anti-static precautions
  - Lack of spare parts
  - Poorly labeled cabling
- Facilities layer
  - Insufficient air-conditioning capacity
  - Power outages
  - Power surges and spikes
  - Fire and flood
  - Physical security
- Egress layer
  - Single power feed from utility
  - Single communications feed
- Personnel
  - Poor quality procedures
  - Lack of discipline
  - Lack of skills
  - Succumb to same disaster as the IT infrastructure
  - Communications unavailable
  - Unable to travel to disaster site

Propose Contingent Solution

Now that the risk factors for the critical business functions within an IT service have been identified and their relative importance and financial implications are clearly understood, the task of crafting a contingency plan can begin.

Service continuity management ensures that the service is available in the event of a service disruption, regardless of the cause of the disruption (disaster, component failure, virus attack, and so on). Service continuance involves two separate but equal functions: failover and restoration.

Failover

Failover is the act of moving the operation of a component from its primary location to a secondary location. This failover may be automatic or manual. For example, a computer has dual redundant power supplies. During normal operation, each supply provides half the power required by the system. When one of the supplies fails, the other automatically begins to supply all the power to the system. This is known as automatic failover.

Another example is a data center site that is destroyed by a tornado. In this case, the entire IT infrastructure must be recreated at a new site located some distance away. This case was formerly known as a disaster recovery. Within MOF, however, this is simply the manual failover of a site.

The need for automated or manual failover to a specific service, application or architecture design can be based on the business impact that occurs if the system was unavailable for a prolonged period of time.

When automated failover solutions are implemented, it is usually in cases where the business impact of an outage is costly in comparison to recovering the systems. That determination can be based on:

- Cost to design an automated, fail-over solution versus the cost of lost revenue and productivity.
- The priority of the services, application, or architectural design that would be affected by a failure.

When the manual failover solution is implemented, it is usually the case that the business impact is not great enough to justify the expense required for an automated failover system.

Restoration

Restoration is the act of bringing the operation of a component from the secondary location back to the primary. This is a very important activity that is often overlooked when creating a contingency plan. People in many organizations have finely detailed plans for how to move the service to a new location in the event of an emergency, but they have no idea how to restore the service when the time comes.

Design for Failover

A countermeasure may be only partially effective by design with, for example, a stand-by system having only half the processing capacity of the primary system. Whether this is acceptable, again, needs to be compared to the definition of the availability agreed upon with the customer to ensure that the countermeasure is considered effective once it is invoked.

Where a particular risk cannot be addressed at an appropriate cost, then either the availability goals need to be re-negotiated with the customer, a strategy of rapid-recovery needs to be adopted instead with its associated impact on availability levels, or the issue passed over to service continuity management and any resulting outage handled as an exception or disaster. Although service continuity management provides contingency plans to handle any disaster, these may involve a prolonged period of downtime before the IT service is fully restored and this factor needs to be taken into consideration during any re-negotiation with the customer.

Such iterative cycles of customer need, risk identification, and design implication may need to be passed back and forth several times during the design and implementation of a highly available IT service.

There are a number of options available to IT professionals to create contingent solutions:

Outsourced Services

Vendor agreements can assist in meeting SLAs on a contingency crisis by ensuring that necessary hardware, people, or recovery locations are available at the moment of need. Considerations on partnering with a vendor, who has a presence within the same geographical region, and other regions, should be evaluated with respect to the findings of the business impact analysis (BIA).

If a natural disaster occurs, the vendor's ability to deliver services or equipment to the affected business may also be reduced in its capabilities; allowing for a failure to meet SLAs. This may come from disruptions in the basic distribution channels of electricity, roads, and airports.

Creating an agreement where a vendor takes over certain aspects of business critical processes in respect to IT transfers the risk to the vendor, who may be better equipped to address the risk. Examples of this form of transfer include ISPs and ASPs hosting solutions. This does not reduce the risk of a disaster occurring, but does place the responsibility and cost of addressing the risk to a specialized industry. As many companies can utilize the same service from the selected vendor, the cost of addressing and mitigating the risk is divided among all the vendor's clients.

Facilities

In some instances, risks to a facility may call for designating a secondary site. This is another area where a vendor may be able to provide such facilities more cost effectively than a company that owns and operates its own facilities. Options for facilities usually fall within the following criteria:

Cold Site, Fixed Center

This solution is sometimes referred to as gradual recovery. This option is applicable to organizations that do not need immediate restoration of business processes and can function for a period of up to 72 hours, or longer, without a re-establishment of full IT facilities. This may include the provision of an empty accommodation fully equipped with power, environmental controls, local network cabling infrastructure, and telecommunications connections, while being available in a disaster situation for an organization to install its own computer equipment.

The accommodation may be provided commercially by a third party, for a fee, or it may be private, (established by the organization itself) and provided as either a fixed or portable service.

A fixed facility may be located at the premises of the third party that provides the service, or specially built at a location owned by the subscriber. There is a need to ensure that all services, including telecommunications, market data feeds, and so on, are established and adequate accommodations are available to house staff involved in the recovery process.

The organization calls on contracts for the supply of required computer equipment including computers, servers, and mini computers. The organization or the contractor (whichever has been formally prearranged) configures the equipment to the organizational requirements and loads all data before a service can be provided.

When choosing a gradual recovery, consideration must be given to equipment that is difficult, if not impossible, to replace if no components are kept by the organization. The same difficulties apply to items supplied by organizations that have since gone out of business, possibly putting the service delivery at risk due to delays or potential problems.

Cold Site, Mobile Center

This option is the same as a cold site, fixed center; except the site is mobile or portable. This site might be erected on a pre-designated location, or near the actual facilities. Considerations for security and needed space should be taken into account, as well as the procurement of needed hardware and possibly labor.

This type of portable facility is typically a prefabricated building provided by a third party and located when needed at a predetermined site agreed upon between the organization and the third party. This may be in a car park or another location some distance from the home site, such as another building.

Warm Site, Fixed Center

This is sometimes referred to as intermediate recovery. This option is selected by organizations that need to recover IT facilities within a predetermined time to prevent impacts to the business process. This typically involves the failover of critical systems and services within a 24- to 48-hour period.

Most common is the use of commercial facilities, which are offered by third party recovery companies to a number of subscribers, thereby spreading the cost across those subscribers. Commercial facilities often include operation, system management, and technical support. The cost varies depending on the facilities requested, such as processors, peripherals, communications, and how quickly the services must be restored (invocation timescale).

The advantage of this service is the customer can have virtually instantaneous access to a site, housed in a secure building, in the event of disaster. It must be understood; however, that the restoration of services at the site may take some time as delays may be encountered while the site is re-configured for the organization that invokes the service and the organization's applications and data needs to be restored from backup storage.

There is a disadvantage in that the site is almost certainly some distance from the home site, which presents a number of logistical problems. The facilities are shared with other organizations, so there can be no guarantee of availability if an incident were to affect two organizations at the same time. However, the recovery companies apply good risk management to the sale of the facilities, which reduces the risk of a multiple invocation. In recent years the number of recovery centers has increased considerably and, together with the falling cost of computer hardware, good deals can often be negotiated for three-, five-, or seven-year contracts.

If the site is needed, there is often a daily fee for use of the service in an emergency; although, this may be offset against additional cost of working insurance. Most commercial agreements limit invocation access to a pre-determined length of time, typically about three months.

It is important that any arrangements of this sort include adequate opportunity for testing at the contingent site.

Warm Site, Mobile Center

Commercial recovery services can be provided in portable form where an agreed upon system is delivered to a customer's site, within a certain time, typically 24 hours. The computer equipment is contained in a trailer and transported to the site by truck. The trailer is outfitted as a computer environment with the necessary services and only needs power and telecommunications links from the site to the trailer for the service to be established. Special measures might need to be taken to make the site secure.

The service provider normally charges an annual fee for such a service, and there is often a call-out charge if the service is invoked. An advantage of this approach is the trailer can be installed close to the main site, subject to the necessary parking consents that are obtained.

Organizations with alternative locations may opt for a mutual fall-back arrangement where accommodation is provided through displacement of non-critical staff at the unaffected building and computer facilities provided via mobile recovery.

Should a mobile solution be utilized or a temporary building be erected on the previous premises, the security of the data as well as the facilities should be taken into consideration.

Hot Site, Fixed Center

This option provides for immediate restoration of services and is usually provided as an extension to the intermediate recovery provided by a third party recovery supplier. An instance where immediate recovery may be required is where the impact of loss of service has an immediate impact on the organization's ability to service its customers, such as a bank. Where there is a need for a fast restoration of a service, it is possible to 'rent' floor space at the recovery site and install servers or systems with application systems and communications already available and data mirrored from the operational servers. In the event of a system failure, users can immediately switch to the backup facility with little or no loss of service.

In the case of a building loss or denial of access, an organization can pay for a limited number of exclusive positions at a recovery center. This is a highly expensive option and is not appropriate for the majority of organizations. If chosen, however, these positions are always available and ready for immediate occupation and use.

Some organizations may identify a need for their own exclusive immediate recovery facilities provided internally. This again is an expensive option but may be justified for a certain business process where non-availability for a short period may cause a very serious impact. The facility needs to be located separately and far enough away from the home site that so it cannot be affected by the same disaster affecting the home site.

For highly-critical, business processes, a mirrored service can be established at an alternative location, which is kept up-to-date with the live service, either by data transfers at regular intervals or by simultaneous update transactions from the live service. Such a service could be used merely as a backup service, but it might also be used for enquiry access (such as reporting) without affecting the live processing performance. This is also useful if there are legal or legislative obligations to safeguard the completeness and integrity of all financial records. As this is essentially spare capacity, under normal circumstances, this spare capacity can be used for development, training or testing, but it can also be made available immediately when a service continuity situation demands it.

Personnel

In a contingency situation, there may be a need to utilize a large number of people to help rebuild the data center in a remote location. The people needed may be for tasks such as unpacking servers or more technical tasks such as configuring routers. If personnel are not available within a company, outside contractors may address this issue.

Design for Recovery

Once the service has been moved to the secondary location, repair of the original system can begin. Once the repair is complete, the service must be recovered to the original site or system. Many contingency plans are developed that outline how to move service to a contingent site, but few have addressed the issue of restoration.

In many cases the recovery plan is simply a variation on the failover plan. The priorities in the plan are probably different. For example, the failover plan may state that high-priority service be moved first. The recovery plan may state that high-cost services are restored first and the high-priority services be restored last. The specific order in which services are restored is dependent on their cost, priority, dependency on other services, and the availability of replacement hardware. Therefore, the restore plan should be a distinct set of documentation, separate from the failover plan.

Designing for Customer Satisfaction During Outages

It is important to note that good customer satisfaction can still be maintained at times of failure despite periods of unscheduled downtime. The key is to establish appropriate and realistic expectations during the requirements definition phase of the availability management life cycle and to articulate a clear understanding of the circumstances under which the service can be expected to fail, and how this is related to the resources being spent to protect it.

Clearly, if the IT service never approaches the levels of availability agreed upon with the customer, then the customer has a right to be dissatisfied. If the reasons for failure and the manner in which the failures are handled fall within expectations, then satisfaction is maintained.

An efficient process for handling and recovering from failures is required; coupled with a good clear communication path to the customer. The customer needs to be kept informed at regular intervals during any recovery processing. Realistic timescales need to be given for when the service is expected to resume.

Formalize Operating Level Agreements

When IT and the customer agree on a cost-effective contingency solution, the agreement needs to be formalized in a document called an operating level agreement (OLA). The OLA serves as one of the building blocks for the service level agreement (SLA) between IT and the customer. The service level agreement is a formal, legal document between IT and its customer. The OLA is an agreement between IT entities.

The operating level agreement includes:

- A definition of the business processing provided.
- Importance to the organization.
- The number of users.
- The business impact of downtime or unavailability.
- The cost of downtime or unavailability and how these costs change over time.
- Hours of service required.
- Critical periods of service, peaks, month-end, deadline processing, and so on.
- Less critical periods of service where downtime is more tolerable.
- Scheduled downtime periods for planned maintenance and upgrades.
- How long downtime can be tolerated before contingency plans need to be invoked.
- Minimum performance characteristics required.

This work is undertaken in close cooperation with the service level management function, as the service level manager is ultimately responsible for the negotiation and documentation of service levels with the customer.

Formalize the Contingency Plan

The contingency plan needs to be a prescriptive guide that can be used by IT personnel to failover and recover the service in the event of a significant event. This document needs to include information on escalation and notification procedures, startup and shut down procedures, communications methods, and status reporting requirements.

Definition of Contingency Levels

The document needs to begin by outlining the various levels of contingency and how to determine if the system is within one of those levels. There are many ways to define these levels, but the following is one that has proven to work in the past.

Table 2. Definition of Contingency Levels

Level	Description	Example
1.	All systems operating properly	Less than or equal to five percent of system users cannot logon after two attempts and Customer search time on phone number takes less than or equal to 40 seconds
2.	Marginal system degradation Ability to meet SLA at risk	More than five percent of system users cannot logon after two attempts or Customer search time on phone number takes greater than 40 seconds
3.	System degradation affecting ability to meet SLA	More than ten percent of system users cannot logon after two attempts or Customer search time on phone number takes greater than 60 seconds
4.	Critical system degradation or More than one system affected	More than 20 percent of system users cannot logon after two attempts or Customer search time on phone number takes greater than 90 seconds
5.	Catastrophic system failure	100 percent of system users cannot logon

There are sophisticated monitoring tools available that analyze a system over time and, using statistical methods, create dynamic targets for each of the previously listed levels. The user of such systems can reduce the number of false alarms presented to the monitoring staff.

Escalation and Notification Procedures

The next section outlines the personnel that must be notified if a contingency level is reached. These personnel need to be chosen based on who can fix the problem and who is affected by the break.

This list may include the relevant times that individuals can be contacted. The following is a sample notification list.

Table 3. Escalation and Notification Matrix

1	2	3	4	5
	Jeff Smith, IT	Jeff Smith, IT	Jeff Smith, IT	Jeff Smith, IT
		Andreas Berglund Finance (7am-7pm)	Andreas Berglund Finance (7am-7pm)	Andreas Berglund, Finance (24 hrs)
			Kim Ralls, Finance (7pm-7am)	Rob Young, CIO (24 hrs)
				Michael Graff, Public Relations (24 hrs)
				Jo Berry, CEO (24 hrs)

Startup and Shutdown Procedures

Startup and shutdown procedures are important for an entire organization's data center. Should a power outage occur, or the need to shut down the entire site arises, the sequence in which it is to be conducted should be readily available. Startup and shutdown procedures can be grouped by technology and by dependency. If a natural catastrophe occurs (such as a flood, fire, or earthquake) then this information is needed to bring up the new data center at another location. Having this information available assists in rebuilding and activating a recovered data center with less time spent troubleshooting dependencies.

Shutdown and startup procedures should be documented so there can be no possibility of misunderstanding the process. The most common method of documenting these procedures is by dependency.

A dependency shutdown document maps out the current network infrastructure with critical divisions in the network, as well as dependencies between the IT layers. This document can contain information that pertains to a router or server within the network that must be online first, after the environment has been shutdown, and then provides information on how to start up the rest of the network or segment in a specific order. This document is scaled in reference to the level of authority on a specific site. For instance, if a remote site has several servers, but is not a main site, the document to reboot the entire site does not include how to reboot the corporate site, but at the corporate site, the start up and shutdown procedures do have the procedures for the remote site.

Ensuring that appropriate permissions are delegated to the appropriate resources to conduct such an activity is also essential to keeping recovery times short. This information is important in developing a recovery process in cases where virus infection, network attacks, or network isolation is needed.

Communications Methods

The contingency plan needs to outline the communications methods that are used by the repair staff to communicate with each other and with the affected process owners. This analysis should include contingency plans for each communication type as well. For example:

- Telephone
    - What if the private branch exchange (PBX) is the affected service? Are handsets available that can be plugged into the incoming legacy telephone service (POTS) lines?
    - What if the PBX trunk lines are unavailable?
- Handheld radios or cellular telephones
    - Can the radios be used in all areas of the building?
    - Is the useful life of the batteries sufficient to outlast the expected length of the service outage?
    - E-mail
    - What if e-mail functionality is the affected service?

Status Reporting Requirements

The contingency plan needs to outline who needs to be notified for each contingency level. It also needs to outline how often this communication needs to take place. For example, a 'level 2' status may dictate that an e-mail be sent to the IT manager. In the case of a 'level 5' service outage; however, the CEO may demand half-hourly updates to address concerns from investors.

Roles and Responsibilities

Principal roles and their associated responsibilities for service continuity management have been defined according to industry best practices. Organizations might need to combine some roles, depending on organizational size, organizational structure, and the underlying service level agreements existing between the IT department and the business it serves.

A small organization may have one person perform several roles, while a large organization may have a team of people for each role (for example, an availability management department). In the latter case, a staff is assigned the job of carrying out the directives they receive from the availability manager.

It is advisable in all cases to have one person who is responsible for implementing the process and can intercede to "make things happen" when the process stalls. This person is the availability manager.

Availability Manager

The availability manager is responsible for managing the activities of both the availability management process and the service continuity management process. This individual is responsible for ensuring that any given IT service delivers the levels of availability agreed to with the customer and for interfacing with all other management processes in pursuit of this goal.

The availability manager:

- Ensures customer requirements are correctly translated into realistic availability goals.
- Ensures appropriate IT budgets are established for protecting IT services.
- Oversees planning activities in relation to designing for availability and recovery.
- Ensures that all risks to availability are identified and appropriately handled.
- Undertakes availability modeling to help select the most appropriate countermeasures, assesses the impact of future changes, and identifies potential improvements.
- Implements cost-effective countermeasures to reduce single points of failure, where possible.
- Ensures that remaining gaps are identified to the customer.
- Ensures that remaining gaps are ultimately handled by service continuity management when required.
- Ensures that the overall IT infrastructure has matured enough to support the availability needs.
- Defines the need for, and helps with, the implementation of availability monitoring processes and tools.
- Ensures availability goals are reflected within appropriate service level agreements both inside and outside the company.
- Manages the day-to-day availability requirements of services.
- Collects and interprets availability metrics on behalf of the customer.
- Forecasts the impact of future availability requirements.
- Participates in the change advisory board to review the availability impact of proposed business and infrastructure changes.
- Manages a continuous availability improvement process.
- Provides consulting expertise for the review and creation of any external contracts that include availability clauses.
- Ensures that contingency plans are in place and up-to-date.

Application Architect

The application architect is responsible for designing applications for IT, especially those that deliver services to IT customers. The application architect should ensure that each application has no single point of failure. This person also ensures that each application responds gracefully to unexpected incidents and does not corrupt data nor create risk to the company through a loss of security. The application architect ensures that applications can be deployed easily in the event that a stand-by site is necessary.

The application architect:

- Designs an application to meet a defined business need.
- Determines the distribution of work in an *n*-tier client-server system.
- Locates network services required for the application.
- Undertakes availability modeling to help select the most appropriate countermeasures, assesses the impact of future changes, and identifies potential improvements.
- Implements cost-effective countermeasures to reduce single points of failure, where possible.

Messaging and Middleware Architect

The messaging and middleware architect is responsible for designing messaging and middleware solutions used by the application architect. This architect helps ensure that each application responds gracefully to unexpected incidents and does not corrupt data or create risk to the company through a loss of security. The architect should ensure that messages can be re-routed easily in the event that a stand-by site is necessary.

The messaging and middleware architect:

- Manages the data communications between system services. This includes managing messaging communications standards such as XML.
- Creates message format standards with internal and external entities.
- Undertakes availability modeling to help select the most appropriate countermeasures, assesses the impact of future changes, and identifies potential improvements.
- Implements cost-effective countermeasures to reduce single points of failure, where possible.

Network Designer

The network designer is responsible for designing network solutions used by the application architect, and the messaging and middleware architect. The network designer helps ensure that the network responds gracefully to unexpected incidents and does not corrupt data or create risk to the company through a loss of security. The network designer should ensure that messages can be re-routed easily in the event that a stand-by site is necessary.

The network designer:

- Creates designs for new network facilities as business needs change and grow.
- Evaluates the design of existing network components as new technologies are developed.
- Facilitates the availability of database, directory, and messaging services.
- Undertakes availability modeling to help select the most appropriate countermeasures, assesses the impact of future changes, and identifies potential improvements.

Hardware Designer

The hardware designer is responsible for designing hardware solutions used by IT. The hardware designer should help ensure that the hardware responds gracefully to unexpected incidents and does not corrupt data nor create risk to the company through a loss of security. The hardware designer should ensure that the system will be easy to repair, even under unusual circumstances.

The hardware designer:

- Determines new hardware needs to meet company objectives.
- Creates initial forecasts of spare parts required.
- Evaluates the existing hardware infrastructure as new technologies are developed.
- Undertakes availability modeling to help select the most appropriate countermeasures, assesses the impact of future changes, and identifies potential improvements.
- Implements cost-effective countermeasures to reduce single points of failure, where possible.

Relationship to Other Processes

Service continuity management is one of the foundational service management functions (SMFs) in the MOF optimizing quadrant. This quadrant seeks to negotiate service level agreements with customers and optimize the IT infrastructure, possibly initiating requests for changes (RFCs) to the IT infrastructure. The list of SMFs in the MOF optimizing quadrant is shown below.

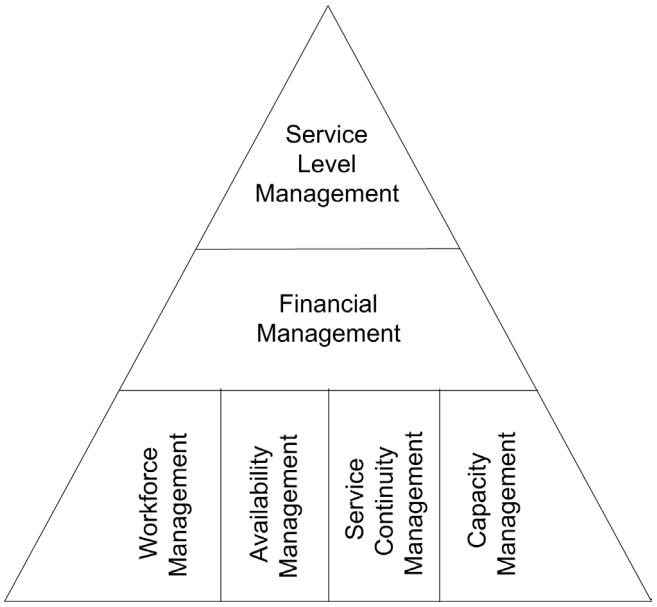

Figure 2
Microsoft Operations Framework Optimizing Quadrant

Service Level Management

Service level management negotiates and manages service level agreements (SLAs) and operating level agreements (OLAs).

Service level management takes primary responsibility for interfacing with the customer and determining which IT services are most crucial to the survival of the company. Availability management draws on this prioritization work and identifies the key IT infrastructure components that support these critical services and determines whether they contain any single points of failure or other risks to availability that can be cost-effectively addressed through the use of appropriate countermeasures. Where no straight-forward countermeasures are available, or where the countermeasure is prohibitively expensive or beyond the scope of a single IT service to justify in its own right, the availability risks are passed to service continuity management to handle.

Financial Management

Financial management acts a filter; ensuring that solutions proposed by availability management, capacity management, or service continuity management can be justified in terms of their cost to implement versus their benefit to the customer. Financial management strives to monitor, control, and, if necessary, recover costs incurred by the IT organization.

Workforce Management

Whenever a new technology is introduced into the IT environment, the people that run that technology must be properly trained and motivated. Workforce management ensures that existing personnel are trained and ready to operate a new availability solution when it is ready.

Contingency plans may involve the restoration of service to a remote location or to the same location after a disaster event. People should not be forgotten as they need safe access to the site, clean water, sewage facilities, and other critical items. If an event has an impact on the IT systems, it is likely to have an impact on the personnel operating those systems as well.

Availability Management

Availability management and service continuity management are closely related as both processes strive to eliminate risks to the availability of IT services. The prime focus of availability management is handling the routine risks to availability that can be expected on a day-to-day basis such as the failure of a hardware component. Service continuity management addresses the more extreme and relatively rare availability risks such as fire or flood.

Where no straight-forward countermeasures are available or where the countermeasure is prohibitively expensive or beyond the scope of a single IT service to justify in its own right, availability management passes these risks to service continuity management to handle.

Capacity Management

Capacity management ensures that appropriate IT resources are available to meet customer requirements by planning for additional resources as current system resource use begins to near the point of full capacity. Service continuity management has a very close tie to this process since the contingency plan may outline reduced capacity capabilities in the event of a disaster. These reduced abilities should be clearly outlined in the OLAs that make up the service.

Contributors

Many of the practices that this document describes are based on years of IT implementation experience by Accenture, Avanade, Microsoft Consulting Services, Fox IT, Hewlett-Packard Company, Lucent Technologies/NetworkCare Professional Services, and Unisys Corporation.

Microsoft gratefully acknowledges the generous assistance of these organizations in providing material for this document.

Program Management Team

William Bagley, Microsoft Corporation

Jeff Yuhas, Microsoft Corporation

Lead Writer

Elias VarVarezis, Unisys Corporation

Contributing Writers

Joe Helm, Unisys Corporation

Vicky Howells, Fox IT

Kristi Moe, Unisys Corporation

Louie Peak, Unisys Corporation

Editor

Michael Laverty, Microsoft Corporation