Overview of the MOF Risk Management Discipline

The Microsoft Operations Framework (MOF) Risk Management Discipline applies proven risk-management techniques to the challenges that operations staff members face every day. There are many models, frameworks, and processes for managing risks-all of which discuss planning for an uncertain future. However, the MOF Risk Management Discipline offers greater value than many others through its key principles, consistent terminology, structured and repeatable six-step process, and a recognition that the MOF Risk Management Discipline needs to be an integral part of the overall operations framework.

Key Principles

An essential aspect of successful IT operations involves managing the risks inherent in running the IT infrastructure. Within the MOF Risk Management Discipline, risk management is the process of identifying, analyzing, and addressing risks proactively. The goal of risk management is to clear the way for the positive impacts (opportunities) of an operations activity while minimizing the negative impacts (losses) associated with that risk. Effective processes for understanding and managing risks will ensure that effective trade-offs are made between risk and opportunity.

In order to implement the MOF Risk Management Discipline, you must have a solid understanding of the following key principles:

Risk Is Inherent in Operations

The only environment that has no risk is one whose future has no uncertainty-where there is no question of whether or when a particular hard disk will fail, no question of whether a Web site's usage will spike or when or how much, and no question of whether or when illness will leave the service desk short-staffed. Such an environment does not exist.

By always keeping in mind that risk is inherent, operations professionals seek ways to continuously make the right trade-off decisions between risk and opportunity and to not become too focused on minimizing risk to the exclusion of all else. IT staff need to stay agile and expect change.

Proactive Risk Management Is Most Effective

Proactive risk management is not achieved by simply reacting to problems. Operations staff should work to identify potential risks in advance and to develop strategies and plans to manage them. Plans should be developed to correct problems if they occur. Anticipating potential problems and having well-formed plans in place shortens the response time in a crisis and can limit or even reverse the damage caused by the occurrence of a problem.

Treat Risk Identification as Positive

Operations staff should always regard risk identification in a positive way; doing so will ensure that people contribute as much information as possible about the risks they face. A negative perception of risk causes people to feel reluctant to communicate risks they perceive. The environment should be such that individuals identifying risks can do so without fear of retribution for honest expression of tentative or controversial views. Managers should support and encourage development of a no-blame environment to foster open communications and promote successful risk management discussions.

Assess Risks Continuously

Many IT professionals misperceive risk management as a necessary, but boring task to be carried out only at the beginning of a project or before the introduction of a new service. Continuing changes in operations environments require process owners to regularly look for new operational risks, reassess

the status of known risks, and reevaluate or update the plans to prevent or respond to problems associated with these risks.
The MOF Risk Management Discipline advocates the use of a structured process that identifies and analyzes risks. This process provides decision makers with information not only on the presence of risks, but the importance, or ranking, of those risks as well.

Integrate Risk Management into Every Role and Function

At a high level, this means that every IT role shares the responsibility for managing risk and that every IT process is designed with risk management in mind. At a more concrete level, it means that every process owner:

- Identifies potential sources of risk.
- Assesses the probability of the risk occurring.
- Plans to minimize the probability.
- Understands the potential impact.
- Plans to minimize the impact.
- Identifies indicators that show the risk is imminent.
- Plans how to react if the risk occurs.

One of the key roles within the MOF Team Model Service Role Cluster might be a service manager. For example, the service manager with overall responsibility for the e-mail service performs all of these tasks to manage the risks that are most important for that service. Other people in that manager's extended staff may perform a subset of those tasks. Everyone will help identify new risks, but perhaps only one or two people will be responsible for estimating probability or making plans to minimize the consequence of that risk.

Shared Responsibility and Clear Accountability

Everyone in IT operations is responsible for actively participating in the risk management process. Process owners are assigned action items that specifically address risks within their service area, and each holds personal responsibility for completing and reporting on these tasks in the same way that they do for other action items related to day-to-day operations. Activities include risk identification within areas of personal expertise or responsibility and extend to include risk analysis, risk planning, and the running of risk control tasks.
Within the MOF Team Model, the Service Role Cluster holds final accountability for organizing risk management activities and ensuring that they are incorporated into the standard processes to meet service level agreements (SLAs).

Use Risk-Based Scheduling

Maintaining an environment often means making changes in a sequence. Where possible, process owners should make the riskiest changes first. The greatest risks tend to be those with the highest level of unknowns. Risk-based scheduling involves making quality trade-off decisions and is important because it minimizes wasted efforts, allowing more reaction time for risk mitigation.

Learn from All Experiences

MOF assumes that focusing on continuous improvement through learning will lead to greater success. Knowledge captured from one experience will decrease the uncertainty surrounding decision making when it is applied by others in later situations. MOF emphasizes the importance of organizational-level

or enterprise-level learning from operational outcomes by incorporating a learning step into the risk management process.

Keep It Simple

Success requires a risk management process that people understand and use. This is a balancing act. If the process has too little structure, people may use it, but the outputs will not be useful; if it is too prescriptive, people probably won't use it at all.

Risk Management Process Overview

The following diagram illustrates the six steps of the risk management process: identify, analyze and prioritize, plan and schedule, track and report, control, and learn. It is important to understand that the process of managing each risk goes through all of these steps at least once and often cycles through numerous times. Also, each risk has its own timeline, so multiple risks might be in each step at any point in time.

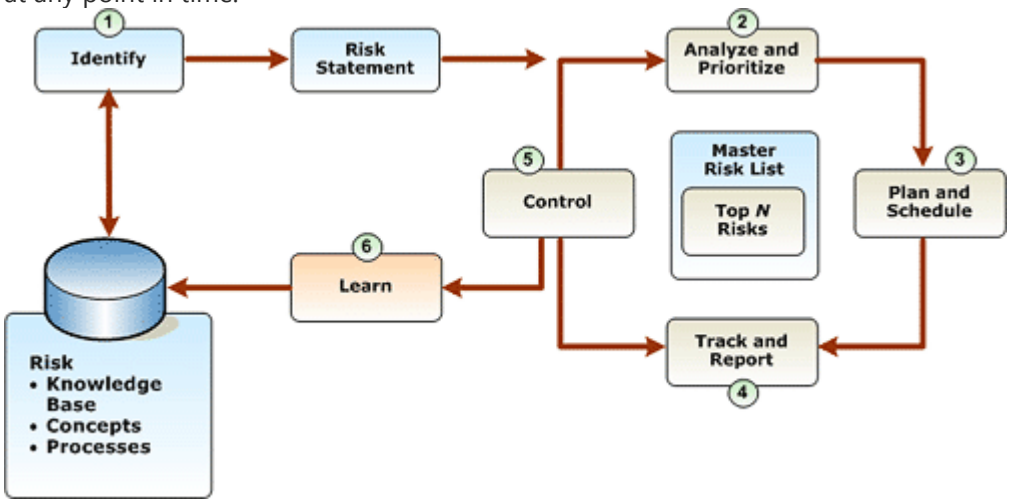

Risk Management Process Steps

The following is a brief introduction to the six steps of the risk management process.

- **Identify** - Risk identification allows individuals to identify risks so that the operations staff becomes aware of potential problems. Not only should risk identification be undertaken as early as possible, but it also should be repeated frequently.

- **Analyze and prioritize** - Risk analysis transforms the estimates or data about specific risks that developed during risk identification into a consistent form that can be used to make decisions around prioritization. Risk prioritization enables operations to commit resources to manage the most important risks.

- **Plan and schedule** - Risk planning takes the information obtained from risk analysis and uses it to formulate strategies, plans, change requests, and actions. Risk scheduling ensures that these plans are approved and then incorporated into the standard day-to-day processes and infrastructure.

- **Track and report** - Risk tracking monitors the status of specific risks and the progress in their respective action plans. Risk tracking also includes monitoring the probability, impact, exposure, and other measures of risk for changes that could alter priority or risk plans and ultimately the availability of the service. Risk reporting ensures that the operations staff,

service manager, and other stakeholders are aware of the status of top risks and the plans to manage them.

- **Control** - Risk control is the process of executing risk action plans and their associated status reporting. Risk control also includes initiating change control requests when changes in risk status or risk plans could affect the availability of the service or service level agreement (SLA).

- **Learn** - Risk learning formalizes the lessons learned and uses tools to capture, categorize, and index that knowledge in a reusable form that can be shared with others.

Risk Lists

The simplest view of the risk management process is that the six steps described previously supply information for a collection of risk lists. These various risk lists can be thought of as a database of risks affecting operations. The concept of a risk database is technology-independent; it could be as crude as a set of index cards, although that would make certain functions (such as sorting, searching, and linking) very labor-intensive and prone to error. The list can be implemented simply as a Microsoft Word document or a Microsoft Excel worksheet, or it can be more effectively implemented using a database application or Microsoft Project.

Note

The size of the risk database is more an indicator of the IT group's thoroughness than an indicator of the health or stability of the IT infrastructure. Using a database application for this purpose should allow you to create customized views or queries into the stored risk information. Four suggested views are: the master risks list, the risks by services list, the top risks list, and the retired risks list. Understanding these views make the six steps for risk management easier to learn and understand.

To access an online example risk list, see Operations Templates.

Master Risks List

The master risks list identifies the condition causing each risk, the potential adverse effect (consequence), outcome (frequently called the downstream effect), and the criterion or information used for ranking, such as probability, impact, and exposure. When sorted by the ranking criterion level (high-to-low), the master risks list provides a basis for assigning priorities in the planning process. During each step in the risk management process, the process owners gather information about operational risks and add that information to the master risks list. It is a regularly updated, or "living," document that forms the basis for the ongoing risk management process and should be kept up-to-date throughout the cycle of risk analysis, planning, and monitoring. Each step in the risk management process builds on the previous step by adding more elements of the risk or draws on the current elements to support decision making. For example, the analyzing step initially adds information about a risk's impact and probability. The process is cyclic, so future passes through the analyzing step may review and revise those impact and probability estimates.

The master risks list is the fundamental document for supporting active or proactive risk management. It enables group decision-making by providing a basis for:

- Assigning priorities.
- Identifying critical actions.
- Highlighting dependencies.

Risks by Services List

The risks by services list is a useful view that allows operations to look at each risk where the consequences of that risk affect a specific business function or service, such as e-mail, customer relationship management (CRM), or payroll. By being able to easily and quickly link risks to their impact on end-to-end services provided by the IT infrastructure, the quality of information, prioritization, and decision making is improved.

Before a risks by services list can be created, it is recommended that IT operations produce a service catalog that lists all of the services currently being provided, a summary of their characteristics, and details of their users and those responsible for their ongoing maintenance.

Top Risks List

Managing risk takes time and effort away from daily operations activities, so it is important for the operations staff to balance the overhead of risk management against the expected savings. This usually means identifying a small number of major risks that are most deserving of limited time and resources. One way to do this is to prioritize the master risks list. The risks at the top of the list, the ones that are important enough to be actively managed, make up a separate top risks list. The size of this list will vary among IT groups, and within one IT group it is likely to vary over time.

Retired Risks List

The master risks list holds all the risks that have been identified, whether or not they are important enough to appear on the top risks list. Some of those risks never go away, such as those related to natural disasters. Others reach a point where they are no longer relevant. For example, the probability of the risk might be reduced to zero, or the source of the risk may leave the environment. Risks specific to an outdated software application are no longer relevant after that application has been completely phased out.

Whenever a risk becomes irrelevant, it is moved from the master risks list to the retired risks list. This list serves as a historical reference from which others can draw on in the future. For example, when risks related to service desk processes are tracked and recorded, and the service desk function is outsourced to another company, some of the service desk risks might be retired. If the service desk function is later brought back in-house, the operations staff can refer to the retired risks list for guidance. Also, people may consult this list as a starting point for identifying new risks.

Finally, if IT operations reduces a risk's probability or impact to zero, then the notes about what was done may benefit other people facing similar risks.

Step 1 - Identifying Risks in Operations

Risk identification is the first step in the proactive risk management process. It provides the opportunities, indicators, and information that allows an organization to raise major risks before they adversely affect operations and hence the business.

This step is closely related to the Information Technology Infrastructure Library (ITIL) term "classification"-formally identifying incidents, problems, and known errors by origin, symptoms, and causes.

Risk Statements

Before a risk can be managed, the operations staff must clearly and consistently express it in the form of a risk statement.

A risk statement is a natural language expression of a causal relationship between a real, existing state of affairs or attribute, and a potential, unrealized second event, state of affairs, or attribute. The first part of the risk statement is called the condition and provides the description of an existing state of affairs or attribute that operations feels may result in a loss or reduction in gain.

The second part of the risk statement is a second natural language statement called the consequence and describes the undesirable attribute or state of affairs. The two statements are linked by a term such as "therefore" or "and as a result" that implies an uncertain (less than 100 percent) but causal relationship. The two-part formulation process for risk statements has the advantage of coupling the risk consequences with observable (and potentially controllable) risk conditions early in the risk identification stage.

Root Cause

When formulating a risk statement, the operations staff should consider the *root cause* or originating source, of the risk condition. Understanding root causes can help to identify additional, related risks. There are four main sources of risk in IT operations:

- **People** - Even if a group's processes and technology are flawless, human actions (whether accidental or deliberate) can put the business at risk.
- **Process** - Flawed or badly documented processes can put the business at risk even if they are followed perfectly.
- **Technology** - The IT staff may precisely follow a perfectly designed process, yet fail to meet business goals because of problems with the hardware, software, and so on.
- **Environment** - Some factors are beyond the IT group's control but can still affect the infrastructure in a way that harms the business. Natural events such as earthquakes and floods fall into this category, as do externally generated, man-made problems, such as civil unrest or changes to government regulations.

These are broad categories, and they can easily overlap. For example, if a newly hired operator undergoes training on the backup software and a week later makes a mistake that causes the backup to fail, is the source of risk "people" or "process?" There are many ways to decide which category a risk fits in, but it is more important to define one way and stick to it, rather than spend time seeking the "perfect" way.

Downstream Effect

The risk identification process results in the identification of the outcome, or downstream effect, of the risk. Understanding downstream effects (total loss or opportunity cost) can help in correctly evaluating the impact that the consequence will have on an organization. There are four main ways in which operational risk consequences can affect the business:

- **Cost** - The infrastructure can work properly, but at too high a cost, causing too little return on investment (ROI).

- **Performance** - The infrastructure can fail to meet users' expectations, either because the expectations were unrealistic, or because the infrastructure performs incorrectly. The reliability of a system can also affect the users' perceptions of the service's performance.

- **Capability** - The infrastructure can fail to provide the platform or the components needed for end-to-end services to function properly or even function at all. For example, consider an enterprise e-mail system that relies upon mail servers, storage servers, gateways or message transfer agents (MTAs), network components, and desktop components. A failure in any one of these components would affect the e-mail service and hence impact the business' capability to communicate effectively.

- **Security** - The infrastructure can harm the business by not providing enough protection for data and resources, or by enforcing so much security that legitimate users cannot access data and resources.

Understanding the characteristics of downstream effect is critical later in the risk identification process when ranking risks to ensure that the most important ones get the attention they deserve since a risk may have a high operational consequence but a low downstream effect, or vice versa.
The following figure schematically depicts the risk identification process along with an example.

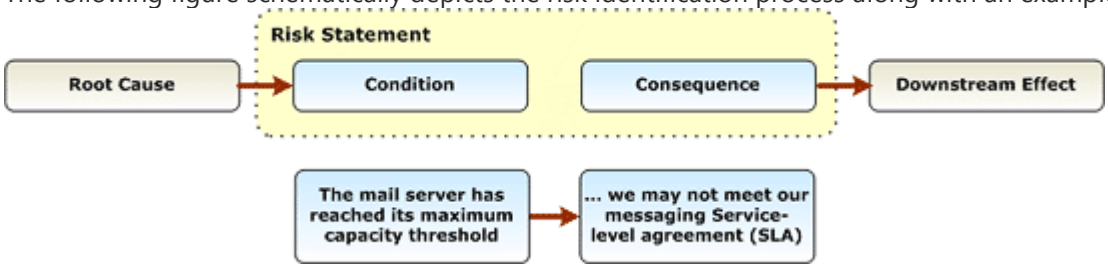

Risks List

The minimum output from risk identification activities is a clear, unambiguous, consensus statement of the risks being faced by the IT operations staff, which is recorded as a risks list. The risk identification step frequently generates a large amount of other useful information, including the identification of root causes and downstream effects, affected service, owner, and so forth.
An example of a risks list produced during the identification step is depicted in the following table. The risks list in tabular form is the main input for the next stage (analysis) of the risk management process and will become the master risks list used during the subsequent management process steps.

Table: Example Risks List

Root cause	Condition	Consequence	Downstream effect
Inadequate staffing	The service desk cannot handle the number of calls it is receiving.	The SLA will not be met and customers will have to wait longer for support.	Reduced customer satisfaction.
Technology change	CRM software vendor plans to withdraw support for the current version of	Existing CRM system will be unsupported.	Reduced sales force capabilities because IT cannot develop the requested enhancements or

	the product.		make any system changes.
New regulatory requirement	All e-mails and attachments need to be stored for eleven years.	Current backup and archiving software cannot accommodate this need.	May result in trading restrictions being imposed and negatively affect the organization's position and image in the market.

Best Practices

These best practices will be beneficial during the risk identification step.

Review Risk Lists and Lessons Learned

A great deal can be learned from reviewing risk databases from similar tasks, talking to process owners about risk management activities in their areas, and reading case studies that identify risks to services or processes. An optimized and mature risk management discipline involves capturing knowledge and best practices from operational activities through the application of such basic knowledge management techniques as consistent taxonomy, risk classification, document management, and advanced search capabilities.

Continual Identification

When a group adopts risk management, the first step is often a brainstorming session to identify risks. Identification does not end with this meeting. Identification happens as often as changes are able to affect the IT infrastructure-which is to say, identification happens every day.

Discussions

Identification discussions are very important. A key to their success is to represent all relevant viewpoints, including stakeholders as well as different segments of the operations staff. This is a powerful way to expose assumptions and differing viewpoints. The ultimate goal of the identification discussion is to improve the organization's risk management capability.

Cause-Effect Matrix

The set of all possible conditions is nearly infinite, and the sheer volume can make it difficult for the operations staff to focus on one at a time, especially during brainstorming. An effective solution, and one that has benefits later in the process, is to subdivide all of the possible conditions into a table with one row for each of the four causes of risk and one column for each of the four types of downstream effect.

		Downstream Effect			
		Cost	Capability	Performance	Security
Root Cause	People				
	Process Technology				
	Environment				

It is now much easier to focus on one cell of the table at a time. For example, IT operations staff might ask themselves, "How might people in the operations group make mistakes that would cause us to do the right work at too high a cost?" Or they might ask, "How could our technology fail to meet customers' performance expectations?" Or more specifically, "How might hardware problems cause the sales group's order entry system to bog down?"

Risk Statement Form

A helpful way to present the information gathered during this step is through a risk statement form, which may add information that will be valuable later during the risk tracking step. In addition to the four parts of the risk statement (root cause, condition, consequence, and downstream effect), a statement form including the following can be very useful:

- **Role or function** - The service management function (SMF) most directly involved with the risk situation.
- **Related service** - Service most affected by the risk.
- **Context** - A paragraph containing additional background information that helps to clarify the risk situation.
- **Related risks and dependencies among risks** - Identify where the consequences of a risk may also be the root cause of or have a direct impact on other risks.

Step 2 - Analyzing and Prioritizing Risks

Risk analysis builds on the risk information generated in the identification step, converting it into decision-making information. In the analyzing step, three more elements are added to the risk's entry on the master risks list: the risk's probability, impact, and exposure. These elements allow operations staff to rank risks, which in turn allows them to direct the most energy into managing the list of top risks.

Risk Probability

Risk probability is a measure of the likelihood that the consequences described in the risk statement will actually occur and is expressed as a numerical value. Risk probability must be greater than zero, or the risk does not pose a threat. Likewise, the probability must be less than 100 percent, or the risk is a certainty-in other words, it is a known problem.
The following table demonstrates an example of a three-value division for probabilities.

Table: Risk Probability Ranges

Probability range	Probability value used for calculations	Natural language expression	Numeric score
1% through 33%	17%	Low	1
34% through 67%	50%	Medium	2
68% through 99%	84%	High	3

Risk Impact

Risk impact is an estimate of the severity of adverse effects, the magnitude of a loss, or the potential opportunity cost should a risk be realized. Risk impact should be a direct measure of the risk consequence as defined in the risk statement. It can either be measured in financial terms or with a subjective measurement scale. If all risk impacts can be expressed in financial terms, use of financial value to quantify the magnitude of loss or opportunity cost has the advantage of being familiar to business sponsors. The financial impact might be long-term costs in operations and support, loss of market share, short-term costs in additional work, or opportunity cost.

The best way to estimate losses is by a numeric scale: the larger the number, the greater the impact to the business. As long as all risks within a master risks list use the same units of measurement, simple prioritization techniques will work. It is helpful to create translation tables to convert specific units such as time or money into values that can be compared to the subjective units used elsewhere in the analysis, as illustrated in the following table. This particular table is a logarithmic transformation where the score is roughly equal to the log10($loss)-1.

High values indicate serious loss. Medium values show partial loss or reduced effectiveness. Low values indicate small or trivial losses. The scoring system for estimating monetary loss should reflect the organization's values and policies. A $10,000 monetary loss that is tolerable for one organization may be unacceptable for another.

Example of a Translation Table

Score	Monetary loss
1	Under $100
2	$100-$1,000
3	$1,000-$10,000
4	$10,000-$100,000
5	$100,000-$1,000,000
6	$1,000,000-$10 million
7	$10 million-$100 million
8	$100 million-$1 billion
9	$1 billion-$10 billion
10	Over $10 billon

When monetary losses cannot be easily calculated, it may be possible to develop alternative scoring scales for impact that capture the appropriate services affected. The following table illustrates a simple example.

Example Alternative Scoring Scale

Score	Criterion	Schedule impact	Technical impact
1	Low	Slip 1 week	Slight effect on performance
2	Medium	Slip 2 weeks	Moderate effect on performance
3	High	Slip 1 month	Severe effect on performance
4	Critical	Slip more than 1 month	Mission cannot be accomplished
100	Catastrophic	Unable to deliver	Mission cannot be accomplished

Risk Exposure

Risk exposure measures the overall threat of the risk, combining the likelihood of actual loss (probability) with the magnitude of the potential loss (impact) into a single numeric value. In the simplest form of quantitative risk analysis, risk exposure is calculated by multiplying risk probability by impact.

Exposure = Probability x Impact

Sometimes a high-probability risk has low impact and can be safely ignored; sometimes a high-impact risk has low probability and can be safely ignored. The risks that have high probability and high impact are the ones most worth managing, and they are the ones that produce the highest exposure values.

When scores are used to quantify probability and impact, it is sometimes convenient to create a matrix that considers the possible combinations of scores and then assigns them to low-risk, medium-risk, and high-risk categories. For the use of a tripartite probability score where 1 is low and 3 is high, the possible results may be expressed in the form of a table where each cell is a possible value for risk exposure. In this arrangement, it is easy to classify risks as low, medium, or high depending on their position within the table. The following table is an example showing probability and impact.

	Low Impact = 1	Medium Impact = 2	High Impact = 3
High Probability =3	3	6	9
Medium Probability = 2	2	4	6
Low Probability = 1	1	2	3

The advantage of this tabular format is that it is easy to understand through its use of colors (red for the high-risk zone in the upper-right corner, green for low risk in the lower-left corner, and yellow for medium risk along the diagonal). It also uses a well-defined terminology: "High risk" is easier to comprehend than "high exposure."

Risk analysis provides a prioritized risk list to guide IT operations in risk planning activities. Within the MOF Risk Management Discipline, this is called the master risks list (described previously in Risk Lists). Detailed risk information including condition, context, root cause, and the metrics used for prioritization (probability, impact, exposure) are often recorded for each risk in the risk statement form.

Best Practices

These best practices will be beneficial during the risk analysis and prioritization step of the risk management process.

Risk Factor Charts

A risk factor chart helps the group quickly determine the exposure it faces for all general categories of risk. One line of such a chart might look like the row in the following table.

Table: Example Risk Factor Chart

Risk	Indicators of High Exposure	Indicators of Medium Exposure	Indicators of Low Exposure
When a hard disk fails, its data cannot be recovered from tape backup.	No one is formally accountable for performing backups. Only one operator has been trained on the new version of the software. The backup operator who has been trained cannot be reached except during his/her shift.	Managers ensure that backups are made every day, but making them is a low-status job assigned to operators with the least seniority. All backup operators attend a one-hour class, but that training covers only the backup software User's Guide and it has no hands-on exercises.	Each week's tapes are sampled and restored to verify integrity. Two backup operators are on shift at all times. Only backup operators who have vendor certification are allowed to make backups without supervision.

Settle Differences of Opinion

It is unlikely that all IT operations staff will agree on risk ranking because staff members with different experiences or viewpoints will rate probability and impact differently. To maintain objectivity in the discussion and to limit arguments, be sure to decide as a group how to resolve these differences before starting this step. Options include a majority-rule vote, picking the worst-case estimate, or siding with the person who has the longest experience dealing with the situation in which the risk event actually occurs.

Measure Financial Impact

It is often helpful to roughly estimate impact in financial terms and record this in addition to the impact's numeric estimate. If several risks have the same exposure value, then the financial estimate can help determine which one is most important. Also, the financial data helps in the planning step to ensure that the cost of preventing a risk is lower than the cost of incurring the consequences.

It might seem that the financial estimate is preferable and could be used in place of a numeric value. In practice, however, financial impact values tend to be a much more labor-intensive way to produce the same top risks list.

If you decide to use a monetary scale for impact, use it for all risks. If a particular risk's impact uses a numeric scale and another's impact uses a monetary scale, then the two cannot be compared to each other, so there is no way to rank one over the other.

Perform a Business Impact Analysis

You should perform a business impact analysis-for example, by using a questionnaire that the users of the service fill out, estimating the importance and impact of service outages. This can help IT understand the service's perceived value, and this might be a factor to consider when ranking risks.

Record the Impact's Classification

Some IT groups find it useful to categorize the nature of the impact, such as security, capital expenditure, legal, labor, and so on.

Step 3 - Planning and Scheduling Risk Actions

Planning and scheduling risk actions is the third step in the risk management process. The planning activities carried out by IT operations translate the prioritized risks list into action plans. Planning involves developing detailed strategies and actions for each of the top risks, prioritizing risk actions, and creating an integrated risk management plan. Scheduling involves the integration of the tasks required to implement the risk action plans into day-to-day operations activities by assigning them to individuals or roles and actively tracking their status.

Planning Activities

When developing plans for reducing risk exposure:

- Focus on high-exposure risks.
- Address the condition to reduce the probability.
- Look for root causes as opposed to symptoms.
- Address the consequences to minimize the impact.
- Determine the root cause, then look for similar situations in other areas that may arise from the same cause.
- Be aware of dependencies and interactions among risks.

During risk action planning, IT operations should consider these six points when formulating risk action plans:

Research

Much of the risk that is present in IT operations is related to the uncertainties surrounding incomplete information. Risks that are related to lack of knowledge may often be resolved or managed most effectively by learning more before proceeding.

Accept

Some risks are such that it is simply not feasible to intervene with effective preventative or corrective measures; IT elects to simply accept the risk in order to realize the opportunity. Acceptance is not a "do-nothing" strategy, and the plan should include development of a documented rationale for accepting the risk but not developing mitigation or contingency plans.
It is prudent to continue monitoring such risks through the IT life cycle in the event that changes occur in probability, impact, or the ability to perform preventative or contingency measures related to this risk. For example, a data center may need to temporarily house servers in a basement room that is at risk of flooding. There may be no alternative location available given the heat and power

requirements. Mitigation or risk transfer would be too expensive and cause too much disruption. In such a case and given the fact that flooding has never occurred before, it may be justifiable to accept the risk and monitor the situation.

Avoid

Risk avoidance prevents IT from taking actions that increase exposure too much to justify the benefit. An example is upgrading a rarely used application on all 50,000 desktops of an enterprise. In most cases, the benefit does not justify the exposure, so IT avoids the risk by not upgrading the application.

Transfer

Whereas the avoidance strategy eliminates a risk, the transference strategy often leaves the risk intact but shifts responsibility for it elsewhere. Examples where risk is transferred include:

- Insurance.
- Using external consultants with greater expertise.
- Purchasing a solution instead of building it.
- Outsourcing services.

Risk transfer does not mean risk elimination. In general, a risk transfer strategy will generate risks that still require proactive management, but reduce the level of risk to an acceptable level. For example, a company with an e-commerce site might outsource credit verification to another company. The risks still exist, but they become the outsource partner's responsibility. However, if the outsource partner is better able to perform credit verification, then transferring the risks can also reduce them.

Mitigation

While the goal of risk avoidance is to evade activities or situations having unacceptable risk, risk mitigation planning involves performing actions and activities ahead of time to either prevent a risk from occurring altogether or to reduce the impact or consequences of its occurring. For example, using redundant network connections to the Internet reduces the probability of losing access by eliminating the single point of failure.

It is vitally important to assign an owner to every mitigation plan, and it is helpful to define the plan's milestones in order to track its progress and its success metrics.

Not every risk has a reasonable and cost-effective mitigation strategy. In cases where a mitigation strategy is not available, it is essential to consider effective contingency planning instead.

Contingency

Risk contingency planning involves creating one or more fallback plans that can be activated in case efforts to prevent the adverse event fail. Contingency plans are necessary for all risks, including those that have mitigation plans. They address what to do if the risk occurs and focus on the consequence and how to minimize its impact. Often IT can establish triggers for the contingency plan based on the type of risk or the type of impact that will be encountered.

Triggers are indicators that tell IT a condition is about to occur, or has occurred, and therefore it is time to put the contingency plan into effect. Ideally, the trigger becomes true before the consequences occur. It may help to think of triggers as warning lights that light up while there is still time to avoid danger. For example, if the condition is that the server runs out of hard disk space, the trigger might be that the server's disk has reached 80 percent of its capacity and is showing an upward trend.

In some cases, the triggers may be date-driven. For example, if the condition is that a newly ordered server might not arrive in time to support the launch of a mission-critical application, a trigger might be set for the latest date on which the server could safely arrive. If the server does not arrive in time

and the trigger becomes true, one contingency plan might be to make use of an existing server from a less-critical service.

Best Practice

This best practice will be beneficial during the risk action planning step.

Prioritize

A *mitigation plan* might have several actions, and the sequence might affect the mitigation's success at reducing, avoiding, or transferring the risk, so it is important to prioritize the steps in this plan.
A *contingency plan* essentially describes how to shift away from normal operations when a condition occurs. Especially if the consequences disrupt many services, it may be valuable to bring some services back online first. Agree beforehand on the order in which to restore service, and decide how long each part can be offline.

Step 4 - Tracking and Reporting Risk

During the risk tracking step, IT operations gathers information about how risks are changing; this information supports the decisions and actions that will be made in the next step (risk control).

Risk Tracking

The risk tracking step monitors three main changes:

- **Trigger values** - If a trigger becomes true, the contingency plan needs to be executed.
- **The risk's condition, consequences, probability, and impact** - If any of these change (or are found to be inaccurate), they need to be reevaluated.
- **The progress of a mitigation plan** - If the plan is behind schedule or is not having the desired effect, it needs to be reevaluated.

This step monitors the above changes on three main time frames:

- **Constant** - Many risks in operations can be monitored constantly or at least many times each day. For example, automated tools can monitor a Web server's bandwidth usage every few seconds.
- **Periodic** - IT operations stakeholders, especially those in the Service Role Cluster, periodically review the top risks list, looking for changes in the major elements. This often happens at staff meetings, change advisory board meetings, OMRs, and so on.
- **As-needed** - In some cases, someone simply notices that part of a risk has changed. This should still be tracked and recorded.

Risk Status Reporting

Risk reporting should operate at two levels-internal and external. For IT operations (internal), regular risk status reports should consider four possible risk management situations for each risk:

- **Resolution** - A risk is resolved, completing the risk action plan.

- **Consistency** - Risk actions are consistent with the risk management plan, in which case the risk plan actions continue as planned.

- **Variance** - Some risk actions are at variance with the risk management plan, in which case corrective measures should be defined and implemented.

- **Changeability** - The situation has changed significantly with respect to one or more risks and will usually involve re-analyzing the risks or re-planning an activity.

Best Practices

The best practices described below will be beneficial during the risk tracking and reporting step.

Review Routinely

Make risk review a part of regular work-for example, making it a permanent agenda item for any recurring meeting. The review can be highly effective without taking very much time. This is the key to managing risks continuously.

Review All Triggers

If the operations staff has highly visible triggers that are automated and constantly monitored, it can be easy to focus on them and overlook triggers that cannot be automated. Forgetting to review such non-monitored triggers means that if one of them has become true, it might not be noticed resulting in further delay of the contingency plan and often compounding the consequences.

Review Trends

Look for trends in risk data. For example, if a particular risk's probability has increased 5 percent every week for the last month, then even though the probability is still low, the trend may justify ranking the risk higher on the top risks list.

Step 5 - Controlling Risk

The fifth step in the Microsoft Operations Framework (MOF) Risk Management Discipline is controlling risk. During this step, individuals carry out activities related to contingency plans because triggers have been reached. Corrective actions are initiated based on risk tracking information.
The MOF Risk Management Discipline relies on existing standard processes and infrastructure to:

- Monitor risk action plans.

- Correct for variations from plans.

- Respond to triggering events.

The results and lessons learned from implementation of contingency plans are then incorporated into a contingency plan status and outcome report so that the information becomes part of the operations risk knowledge base. It is important to capture as much information as possible about problems that occur or about a contingency plan when it is invoked to determine the efficacy of such a plan or strategy for risk control.
At first this step may not seem necessary, and the distinction between it and the tracking step may be unclear. In practice, the need to act is often detected by a tool or by people who don't have the

required responsibility, authority, or expertise to react on their own. The controlling risk step ensures that the right people act at the right time.

Best Practices

The best practices described below will be beneficial during the risk controlling step.

Communication

The risk controlling step relies heavily on effective communication, both to receive notification that parts of risks and plans have changed, and to ensure that the right people take action at the right time. The risk controlling step can't be effective unless communication within IT is also effective.

Step 6 - Learning from Risk

Learning from risk is the sixth and last step in the Microsoft Operations Framework (MOF) Risk Management Discipline and adds a strategic, enterprise, or organizational perspective to risk management activities. Risk learning should be a continuous activity throughout the entire risk management process and may begin at any time. It focuses on three key objectives:

- Providing quality assurance on the current risk management activities so that the IT operations group can gain regular feedback.
- Capturing knowledge and best practices, especially around risk identification and successful mitigation strategies-this contributes to the risk knowledge base.
- Improving the risk management process by capturing feedback from the organization.

Capturing Lessons About Risk

Risk classification is a powerful means for ensuring that lessons learned from previous experience are made available to the groups performing future risk assessments. The following two key aspects of learning are often recorded using risk classifications:

- **New risks** - If IT operations encounters an issue that had not been identified earlier as a risk, it should review whether any signs (leading indicators) could have helped to predict the risk. You may need to update the existing risk lists to help identify risks in the future. Alternatively, you might have identified a new operational risk that should be added to the existing risk knowledge base.
- **Mitigation strategies** - The other key learning point is to capture experiences of strategies that have been used successfully (or even unsuccessfully) to mitigate risks. Use of a standard risk classification provides a meaningful way to group related risks so that operations can easily find details of risk management strategies that have been successful in the past.

Best Practices

The best practices described below will be beneficial during the learning from risk step.

Risk Review Meetings

The risk review process should be well managed to ensure all learning is captured. Operations management reviews (OMRs) as well as specific risk review meetings provide a forum for learning from risk. They should be held on a regular basis and, like other reviews, will benefit from advance

planning, development of a clear, published agenda, participation by all participants, and free, honest communication in a "blame-free" environment.

Risk Knowledge Base

The risk knowledge base is a formal or informal mechanism by which an organization captures learning to assist in future risk management. Without some form of knowledge base, an organization may have difficulty adopting a proactive approach to risk management. The risk knowledge base differs from the risk management database, which stores and tracks individual risk items, plans, and status for a specific service.