ITSM.express

Version 1.1, 18 April 2024

The Team

Project Manager

Štefan Ondek

Authors

Johann Botha
Etienne Shardlow
Dolf van der Haven

Reviewer

Suzanne van Hove
Wim Hoving

Disclaimer

We recognize the many views of service management – is it "IT service management" or just "service management"? Or something else? To answer that question, in terms of this document, we have to ask, "What service doesn't sit on some form of technology?" The answer to that question, is simple – there really isn't a service today that doesn't utilize some form of technology, digital information, etc. Therefore, as our name suggests, we are talking about IT service management, but we all recognize that the information within this document is applicable outside of the technology/digital vertical.

Therefore, our definition of service management is the whole of an organization's capabilities, processes and structure that supports its staff's activities to deliver services throughout their lifecycle to provide value to its customers and the organization itself.

IT service management (ITSM) is a subset of service management, focused on services that primarily consist of information technology elements.

For ease of reading (and typing), we will use the term "service management" and our examples/discussions will focus on the technical enterprise.

Preface

This is a **minimalist, free manual covering the essentials of service management**. It has been written specifically so that it can be easily learned, applied and taught. The audience is primarily people and organizations that are starting their journey in service management, building a new service management system or improving an existing one. There is no reference to existing commercial frameworks – the only external references that are made are to standards. This is done so that this ITSM.express guidance remains free and accessible to everyone who is interested in service management.

The content, while minimalist, remains consistent with ISO/IEC 20000-1:2018. While ISO/IEC 20000-1 and most other (IT) service management systems treat service management actions independently, this manual shows how to create an effective management system through four major service management actions: Define, Produce, Provide and Respond.

This manual has been **written by veterans and top names of this industry**: Johann Botha, Dolf van der Haven, Etienne Shardlow and reviewed by Suzanne Van Hove and Wim Hoving. The editor and "mastermind" behind it is Stefan Ondek, himself a veteran trainer, consultant and subject matter expert in the areas of project and service management. These individuals performed all the work on this manual as volunteers without any compensation. Big thanks to them.

No guarantees are provided for any sort of use or application of this manual. Use your own judgement.

Table of Contents

Define

The Define stage covers the design of a service, including basic organizational processes such as governance and risk management. Here, the foundations for a management system supporting the service lifecycle are laid.

Governance

Objectives

Governance is a function that provides guidance to the organization in the form of direction, policies and general oversight of the overall functioning of the organization. It keeps in mind the goals of the organization in terms of business outcomes, such as profitability, customer and employee experience and compliance to applicable laws and regulations.

According to standards such as ISO 37000 (Governance of organizations) and ISO/IEC 38500 (Governance of IT), the main group in an organization that performs these functions is a "governing body", for example a board of directors or an otherwise relatively independent person or group at the top of the organization. In practice, however, governance can and should be performed at every layer of the organization by appropriate leadership for their scope of responsibility.

Benefits

Good governance is crucial for running any organization, including service providers. The benefits are as follows:

- Better oversight of the performance of the organization as a whole;
- Accountability at all levels of the organization for its operational and strategic performance;
- Clarity on the business direction of the organization;
- Better awareness among employees of the relevance of their activities in the greater context of the organization's purpose;
- Improved decision-making based on the observation of relevant information.

Key Concepts

Good governance is based on a number of principles that every organization can set for themselves, but are likely similar to one or more of the following:

- Purpose of the organization;
- Generation of value for the organization and for its consumers;
- Establishment of a strategy for the organization
- Providing oversight of the organization's operations;
- Engaging stakeholders, such as consumers, employees, authorities;
- Governance of risk;
- Focus on sustainability.

Process

The principles mentioned above and possibly others provide the starting point for building the governance function. Practical activities based on these principles are the following:

- Determine the mission and vision of the organization: what is it that the organization wants to achieve? How do you want to do this?
- What are the factors that influence the activities of the organization, both internally and externally? Think of consumer requirements, legislation and regulation, sustainability goals, availability of resources and overall strengths and weaknesses of the organization at this point in time.
- Who are the stakeholders in the activities of the organization? There are the employees and managers of the organization itself, but also consumers, authorities, perhaps even media and shareholders who are interested in how the organization performs and what the level of success of it is.

There is a simple model in ISO/IEC 38500:2024 that shows what activities are to be performed as part of a governance function. These are as follows, abbreviated as EDMS:

- **Evaluate**: assess how the organization is doing, based on the principles mentioned in section 1.2.
- **Direct**: provide direction to the organization to change its activities if this is deemed necessary based on the evaluation done in the previous step.
- **Monitor**: monitor how the organization performs based on indicators such as key performance indicators (KPIs), key risk indicators (KRIs) or others.
- **Stakeholder Engagement**: provide ample communication about what the governance function wants to various internal and external stakeholders, so that the organization below it knows what it is supposed to do operationally.

These four activities are very easy to implement as responsibilities of all levels of management in the organization. A central governance function such as a board of directors should do all four activities for the organization as a whole, so that managers below it have the necessary guidance to implement the direction received in the operations of the organization. But a leader of a team of employees should do the same at their own level: evaluate the performance of the team, provide direction to the team on what needs to be done to improve its goals, monitor how the team is doing based on established KPIs and communicate frequently about the direction given and the performance of the team.

Managers in the organization are therefore wearing two hats: on the one hand, they need to put into practice the (governance) direction they receive from higher up in the organization. On the other hand, they need to provide governance direction themselves to the team or part of the organization that reports to them. This split between governance (the EDMS activities) and management (translating governance direction

into operational practice) is sometimes exaggerated but in reality, governance and management are an integrated whole for each management level in the organization.

In practice, there is even a third “hat” that managers wear, which is the one that requires them to manage upwards: it is the “monitor” function of governance that provides insight to higher management, based on which the latter can evaluate the performance of the organization below them and provide new direction as required.

Common Pitfalls

Common issues with governance include the following:

- Governance is not government – government is a political entity at various levels in society to guide the people within their remit and set political direction. Governance, even though it has similar activities associated with it, is a function at various levels in a company or other organization, but has no political interests,
- Governance is confined to a Governance Board at the top of the organization only. It should really be performed at every level in the organization.
- Governance and Management activities are not clearly separated. There should be a clear distinction between the strategic activities of Governance (Evaluate,-Direct,-Monitor,-Communicate) and the Management activities that are required to put the received direction into practice in the organization.
- Governance is misinterpreted as a schedule of meetings between various stakeholders. Governance has clear objectives and activities associated with it, as described above. It obviously requires certain meetings, but the purpose of those must be clear to the participants having Governance responsibilities.

Further Reading

Refer to the ITSM.express website for additional information on input to governance and governance deliverables.

Risk Management

Objectives

A risk can be defined as any uncertainty that may impact the objectives of the organization. These uncertainties may be either positive or negative. A positive risk is something that may stimulate attainment of the objectives and is therefore also referred to as an “opportunity” that the organization may want to pursue. A negative risk is how people usually consider risks: an impediment to achieving the organization’s objectives that should be avoided or at least minimized. Both types of risks are relevant for risk management. Risks are an important part of the input to the governance function at all levels of the organization. Risk management should therefore be performed at all levels of the organization as well. Risk management is part of many other areas of service

management, including Information security management, capacity management and other processes.

The default standard for risk management is ISO 31000 (Risk Management), which this section is generally aligned with.

Benefits

Similar to governance, risk management is a foundational practice of running an organization. Benefits of performing risk management are the following:

- Early warning of possible threats and vulnerabilities;
- The possibility to intervene if risks materialize by putting effective controls into place;
- Clear oversight of the threat landscape;
- Awareness in the whole organization of what may be in the way of generating its business outcomes.

Key Concepts

Risks are events that may impact the organization in various ways: there may be risks related to keeping the information about the organization or its customers secure (information security risks), risks to the operational performance of the organization (the ability to meet its KPIs), risks of not meeting external regulations, such as privacy laws or other regulations (compliance risks), or risks related to external suppliers (third party risk). In risk management, anything goes, really, but it needs to be determined how serious a risk is in reality.

Risk are usually classified based on two of its aspects: the *likelihood* that the risk actually occurs in practice (e.g. on a scale of 1-5) and, if it occurs, the *impact* it would have on the organization (e.g. on a scale of 1-5). Once you determine likelihood and impact, you can calculate a *risk level* by taking the product of the two:

Risk Level = Likelihood x Impact

This risk level serves as an overall indicator of the severity of the risk.

Process

Risk management as defined in ISO 31000 has several phases:

1. Risk Assessment is the overall stage where the following three process steps are performed:
    a. Risk Identification: based on the scope of the risk management process, the context of the organization and the criteria for what is deemed acceptable, risks are identified. At this stage, identification consists of determine threats that may affect vulnerabilities in the organization.

b. Risk Analysis: the threats and vulnerabilities in the previous step are analysed by people in the organization that have the ability to determine the potential impact of the risk if it materializes and the likelihood of the risk actually materializing. This leads to the risk level indicated above.
c. Risk Evaluation: Based on the risk level that was determined during the risk assessment, it should be determined what action needs to be taken.

Actions for risk mitigation depend very much on the nature of the risk, but fall in three main categories:

1. **Treat the risk**: in this case, you decide to take action to reduce either the likelihood of the risk, or the impact of it, or both. The action taken is known as putting a *control* into place. The control is an administrative (write and implement a process or policy), technical (implement a firewall) or physical (add locks to doors) measure that reduces the likelihood or impact of the risk.
2. **Accept the risk**: it may be decided that the risk level is not high enough to warrant spending time or money treating the risk. In other words, the risk is at an acceptable level. This is a decision that should not be taken lightly: only the appropriate level of management should be able to accept risks for their own scope of responsibility and only up to a certain risk level. For instance, it can be a policy that only risk up to level Medium can be accepted and then only by the level of management that is responsible for the part of the organization that is impacted by the risk. This acceptable risk level is known as *Risk Appetite*.
3. **Transfer the risk**: If the risk is of a nature that another team, another part of the organization or even an external party is the one that is responsible for taking action to handle the risk, the risk may be transferred to them. Transfer involves contacting the other party, explaining the risk and your assessment of it to them and getting them to acknowledge that they are responsible for handling the risk. The other party can then decide themselves if they think it is necessary to treat the risk or not. In any case, if the risk actually occurs, it is the other party that is accountable for the impact of the event on the organization. For example, if you identify a risk that unauthorized staff may be able to enter a restricted area in a building, such as a data centre, you may want to transfer this risk to a group in the organization that is responsible for the physical security of the facilities. That physical security groups should then acknowledge ownership of the risk and determine if they need to take action to secure the restricted area.
4. **Pursue the risk**: In the case of positive risks (opportunities), the organization may want to pursue it, meaning that the occurrence of the risk should be stimulated to obtain the positive impact of the risk. This can be done, for example, when it is expected that financial conditions, such as tax regulations or energy prices, are expected to improve, making the provision of services cheaper and therefore improving the service provider's market position.

If a risk gets treated, meaning some form of control is put into place, the likelihood and/or impact of the risk should be lower than the original likelihood and/or impact. This results in a new risk level, again calculated as Likelihood x Impact, that is known as the

Residual Risk. The residual risk is whatever risk is left after putting a control into place. A control often only reduces the risk level down to a certain point but doesn't eradicate the risk completely. The residual risk is an indicator of what risk remains after treating it and should be below the acceptable risk level (i.e. the Risk Appetite). If it is still higher than the risk appetite, additional controls may have to be put into place to lower the risk level sufficiently.

Common Pitfalls

- Risk management is neglected. By putting a practical and easy risk management framework into place, risk management doesn't need to be cumbersome, but can be integrated in the day-to-day operations of the organization.
- Too many risks are accepted rather than treated. This is a sign of lazy risk owners – rather than analysing risks and putting controls into place they take the risk and ignore the possible impact on the organization's business outcomes.
- Controls don't get periodically checked for their effectiveness. When a control to treat a risk is put into place, its effectiveness needs to be checked periodically, because the threats may change and eventually break through the controls. Controls may have to be hardened or replaced by other controls in order to remain effective.
- The person flagging the risk automatically owns it. Risk ownership should be with the person in the best position to deal with it, which is not necessarily the person who flagged it first. A "touch it, own it" attitude in the organization will discourage people from flagging risks in the future.

Further reading

Refer to the ITSM.express website for further details on risk management concepts and practice.

Consumer Interaction

Objectives

A service is developed in order to be used by its end-users or consumers. It is possible to develop a service without interacting with consumers, but this will create a service that may not be interesting for them and therefore doesn't get used. Interaction with consumers needs to take place from the earliest moment of developing a service until the moment it gets decommissioned.

Benefits

Benefits of consumer interaction are the following:

- Better understanding of what the consumers of the service need in order to meet their own business outcomes;
- Regular verification of consumer experience and satisfaction with the service, so that the service can be improved if needed;

- Regular reporting to consumers about service performance so they can verify if the service performs as expected and agreed;
- Guidance for service development to make sure the service meets the needs of the consumers.

Key Concepts

Consumers are rarely interested in your service because of the service itself: they usually want to use the service to achieve their own goals. Somehow, consumers derive *value* from using your service to achieve their own outcomes. This consumer value can be something entirely different from the value that you as the service provider derive from the service: your value may be obtaining revenue, getting happy customers, gaining market share or supporting your wider company's business outcomes. The consumer of the service usually sees the service as an *enabler* for their own outcomes, such as being able to send and receive emails, process data, have financial transactions handled.

The first step in interacting with the (potential or existing) consumers of your service is to determine what *they* believe is the value of the service. This is a very basic question, but it is also the most important one if you want to be able to get consumers to use your service. From the value statement the requirements for the service follow, as well as the basis of forming relationships with the consumers.

Process

From this consumer value definition, you can continue deriving the requirements they have for the service you provide. What aspects of an email service do they need to do their jobs most efficiently (for example shared mailboxes, scheduled sending, etc.)? What aspects of a financial service do they need to achieve their business outcomes (for instance payroll, accounting, etc.)?

Different consumers will have different requirements. It makes sense to collect a list of actual and potential requirements for your service and prioritize those based on what the majority of the consumers need and what is achievable in a certain timeline. There will be a set of common service requirements among a majority of consumers that should be considered part of a basic service. Then there are requirements that stretch the basic service and add functionality that benefits a large group of consumers: these can be considered "premium" functionality at first or be put on a road map for development in the near future. A third category of requirements are the completely custom requirements that only apply to a single or a handful of consumers. For these, you need to consider whether it is worth developing a fully custom service for just a few consumers or not. It depends on your own business model if you want to focus primarily on standard services or have the flexibility to customize them.

Consumer requirements, as well as your own, internally gathered requirements feed into the service development process. This process is discussed later on in the Produce section of this book.

Relationships with the consumers of your service don't end once you have gathered the requirements and sold or provided the service to them. Consumers need to be involved every step of the way so they can voice their feedback about the service you provide to them. This is known as Consumer Experience (CX) management and, at an end-user level, User Experience (UX) management. CX/UX cover the whole of how the users of your service feel about it. This includes usability, functionality, their experience when acquiring the service, billing (if applicable), after-care, requests for change, resolution of incidents and any improvements you make for them. Every time a consumer interacts with your organization or your service, there is a touchpoint and every touchpoint results in the consumer having a certain experience. The core of consumer relationship management is that you can determine that customer experience at those touchpoints, take it on board and, if necessary, implement improvements if the customer experience tells you so.

While keeping an eye on the customer needs, don't forget about the needs of your employees: employee experience is at the foundation of providing successful services. It is the employees that do all the work for the customers and it is therefore important to make sure they feel happy doing their jobs. Employee experience can be measured in similar ways as customer experience, but doing regular surveys or by getting feedback in individual or team-level discussions. This feedback should then be taken to the Continual improvement process to make changes to the services where needed.

A final aspect of interacting with consumers is to provide them with reporting about the service they take from you. This can have various forms, but often relies on certain service targets you agree with the consumers. Service targets can be similar to the following:

- Timeliness of response to service requests, incidents, change requests;
- Availability of the service;
- Performance of the service;
- Reliability of the service;
- Consumer experience (including ease of use, effective end user support, fit-for-purpose service features);
- Accuracy of the service; etc.

Service targets can simply be part of a default contract or be customized for specific consumers. In either case, they would be referred to a Service Level Agreement (SLA) that, if focused more on CX/UX, can include an Experience Level Agreement (XLA).

Common Pitfalls

- On-size-fits-all attitude. This means that it is expected that all consumers will be equally happy with the service, no matter what their own needs are. In practice, every consumer has a different opinion about what the service should do for them, so customization may be needed to make the service fit their individual needs.
- Consumer interaction does not happen frequently enough. It depends on the type of consumer if more or less regular interaction is needed with them. Some consumers prefer monthly reporting, others are fine with an annual performance report. Customization of consumer interaction may be needed to tune its frequency to the needs of individual consumers.
- Interaction happens the same way for all consumers. Similar to frequency, the type of consumer interaction may need customization in order to meet their specific needs and expectations.
- Internal targets are not aligned with consumer SLAs. This is unfortunately quite common: key performance indicators for the teams supporting the service may be misaligned with what has been agreed in consumer SLAs. For example, internal response time targets may make it impossible to meet a response time SLA with a consumer. Internal and external targets need to be aligned and staff supporting the services need to be aware of both their own KPIs and of consumer SLAs.

Further Reading

Refer to the ITSM.express website for further details about customer interaction, including measuring customer experience and reporting on it.

Produce

The Produce stage of service management develops the service by building, implementing and testing it. Here, processes such as change management and release management are used to control the provision of the service in the live environment before consumers will use it.

A service is a way to deliver value to customers by fulfilling the customers' defined needs. Usually, a service is therefore intangible, although it can also be delivered together with a (tangible) product.

Build

Before services can be delivered, they need to be created, meaning that all service components are created and gathered, the organization is made ready to support the service and the service is built based on the service design. Ideally, this is done in collaboration with the consumer, so that both the service provider and the consumer get the most value out of the creation of the service. The value derived from the service may be very different for the service provider and the consumer, though, but both sides need to be considered when building the service. Consumer value can, for example, be facilitating a business process that leads to fulfilling one of their business outcomes; service provider value can be gaining or retaining market share, increasing customer satisfaction or revenue generation. Service features should be based on the customer needs they are to satisfy, but that is not enough; other input should also be considered:

- The Governance model and strategy of the organization (is this service we should deliver?);
- The organization's capabilities, or the capabilities of partners (can we deliver this service?);
- The approved organizational architecture including the technical architecture used by the organization (can we effectively build and maintain this service?).

When thinking about customer needs the organization need to consider not only functionality, but also the warranty aspects of the service they are about to design or build and how they will approach adding the service and once added to its portfolio, offering and maintaining the service.

Service Design

Conditions of use

Service providers must understand how and where services will be used by customers as these determine the warranty design of a service.

Warranty elements to consider are as a minimum:

- **Service Demand**, which determines the load placed on resources used to deliver the service (for instance, how many sites, how many users, how many transactions) and should be used in architecting and designing the service so that not only the features customers need will be present, but also the capacity to always deliver value to customers.
- **Service Availability**, which considers how the customer will use the service, under what conditions and what customers considers acceptable level of availability. When should the service be available, where should it be available, what is considered as available (for instance, slow response may be considered as unavailability, so for a computerised system, this will include the elements like transaction response time and transaction the volume the service should be capable of handling.)
- **Service Capacity**, which is a translation of the above two elements into the technical capabilities of the underlying building blocks of the service. Once again, in an computerised system at the lowest level this translates into elements such as network bandwidth, database volume and read/write speed, storage capacity, processing power (usually CPU and memory capacity.)

Needs Analysis

In-depth customer needs analysis should be conducted. This should be viewed in context of the competitive environment to determine if the new (planned) service can be viably delivered to customers.

Needs analysis should not only consider the features customers desire (the utility the service should provide to customers,) but also the conditions of use to be able to determine warranty elements of the service.

It is best if needs analysis is based on a clearly defined analysis method that not only depend on customer or user interviews but also observation of work done, understanding the needs of similar organizations and what's on offer in the market (the competitive landscape).

Designing and building a service with no unique or distinct value proposition seems arbitrary and organizations should aim to offer more value to customers or better value to customers.

Understanding how and where the service will be used forms a key part of analysis and often this aspect needs to be translated into more technical requirements by specialist.

Build or Buy

Once the Service Design is completed, organizations critically need to look at their capability to deliver on all aspects of the service once build. Often this means that organizations choose to for strategic partnerships for the delivery of some of the aspects of the service.

From an architecture perspective, organizations may also decide to use one of three options when thinking about the components of the service:

1. Develop the components themselves and build it from scratch;
2. Develop some components and use existing building blocks to complete the offering;
3. Buy or acquire standard building blocks which can be assembled in a unique configuration and only need configuration (rather than development).

The cost of these options can vary widely, and the organization should not only consider their capabilities, but also build a financial model to determine the short- and long-term viability of the design.

Services that involve Multiple Vendors

If an organization delivers a service that depends on other service providers, part of the design criteria should be to ensure that other parties can fulfil their commitments as the overall performance of the service will depend on all its components.

Service level agreements must be put in place and means devised on how these will be managed, measured and what actions should be taken when service quality is threatened. Most people immediately think of the legal ramifications but that may already be too little too late. What actions should be taken by all parties to ensure service degradation is actioned and controlled to minimize the negative effect of service degradation on customers.

Approaches like Service Integration and Management (SIAM) provide detailed insight into this service outcome.

Service Monitoring

If possible, monitoring systems should be designed as part of service design that will enable the service provider to proactively action service degradation.

We will deal with this aspect in greater detail in the service monitoring section.

Service Aggregation

Service providers often also manage services offered to the client by other vendors. This is not the same situation as described above, but rather service providers contracted by the customer directly to deliver services, that often may impact on service delivery as a whole.

Often in scenarios like this, service providers do finger-pointing when an outage or service degradation occurs, rather than work together for the good of the customer. It may be prudent for the customer to appoint a service agitator that would manage all service providers end-to-end.

Approaches like Service Integration and Management (SIAM) provide detailed insight into this service outcome.

Change Management

Objectives

The purpose of change management is to control the changes implemented in a production environment. The word "control" may sound restrictive but does not need to be so in practice. There is often a balance to be found between the ability to implement service improvements quickly and flexibly and the need to keep the services stable. It is this balance that change management should provide.

Benefits

There are various benefits of having a good change management process:

- Ensuring the stability of services when changes are deployed;
- Ensuring a minimum delay in deployment of changes;
- Informing all relevant stakeholders about planned changes;
- Facilitating prioritization of less service impacting change types.

Key Concepts

Change management starts with a Request for Change (RFC): this is a request coming from an end-user, a development team or another party to change something in the service you provide. RFCs can involve very minor changes, such as modifying a network interface's speed, or much larger ones, such as upgrading the hardware of a server.

A change management policy should be put into place to guide how various types of changes should be handled:

- Large, service impacting changes may have to go through a project that multiple stakeholders work in to implement the change without disruption. This project would still include the steps required by the change management process;
- Minor changes may be pre-approved and implemented without further delay. These would often be sent into the Request Fulfilment process, which is in fact a shortened version of a change management process (see the Respond section for details);
- All other changes would go through the regular change management process.

There is a category of change that is part of resolution of an incident and is called an Emergency Change. Emergency changes need to implemented without any delay in case there is an urgent need to fix an incident, such as a Distributed Denial of Service (DDoS) attack or a serious software bug that needs patching immediately. A separate

procedure may be in place on how to handle these emergency changes. This procedure would usually initially bypass the regular change management process as outlined below and deploy the change immediately. Afterwards, the regular steps would still need to be followed to document what has been done.

As described, a common change management process can govern all types of changes.

Process

When the service provider receives an RFC, the first step after registering it is to let it go through *Triage*, meaning that the nature of the change should be assessed so that its potential impact can be determined. The result of triage is to send the RFC either into the regular change management process, or to Request Fulfilment or to a project organization to handle.

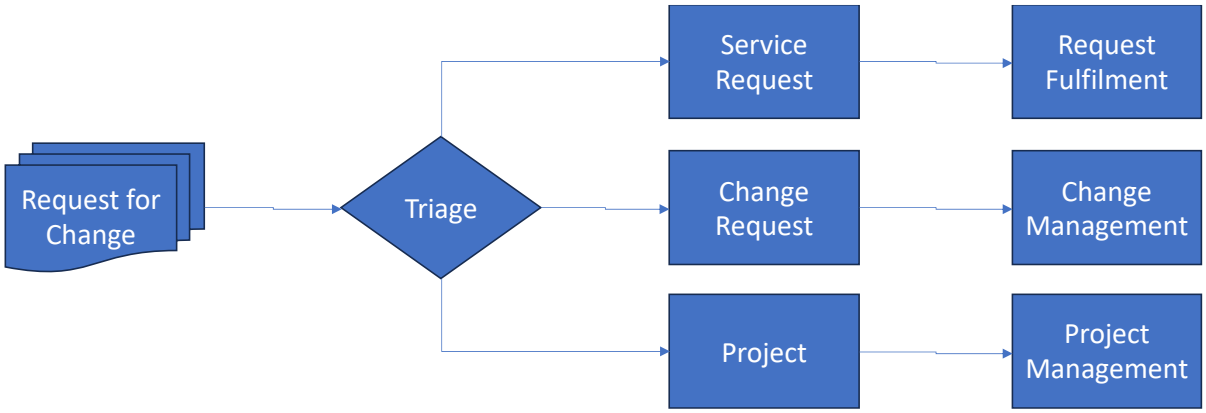

If an RFC goes into the regular change management process, it should be evaluated for its validity by the team that is supposed to implement it. The team may have to consult the requestor to verify details and agree a time period in which the change can be deployed (the Change Window). A roll-back or back-out plan may have to be created to cater for the situation that the deployment fails and the change needs to be rolled back.

In large and complex environments, it may be necessary to have a separate group oversee all RFCs that are received and verify if some of them may conflict with each other and should therefore be deployed separately. This group is referred to as a Change Advisory Board (CAB) and, as the name suggests, advises how RFCs need to be handled. Involvement of a CAB should not lead to delays in deployment, though, but may be needed to ensure stability of the services and prevent change-related incidents.

The approval of a change should be done by the team that has most knowledge about the nature of the change and its impact on the live services. It is therefore done at the lowest possible level in the organization that has the authority to approve change requests.

Once approved and scheduled, the RFC goes into the Release and Deployment processes (see the next section).

In fully automated environments, such as DevOps development, all of the above steps may be automated as well, as long as the right criteria are set for change triage, assessment and approval. This ensures a minimal delay in deployment of changes.

Common Pitfalls

There has been a lot of criticism on change management processes as implemented by various organizations, usually quoting unacceptable delays in getting RFCs deployed. The following pitfalls are at the core if this:

- All changes must be run through a CAB. This is wholly unnecessary as it always introduces delays for changes with a low impact. Only use a CAB in complex environments where there is a risk of conflicts between different RFCs and make sure this CAB meets frequently.
- Many approvals need to be obtained before a change can be deployed. As long as the relevant stakeholders are informed, only one real approval is needed for a change, which is the one from the team actually deploying it. To reduce the need for approvals, also see if simple changes can be performed without having to go through an approval process: this concept is known as “standard changes” that can be implemented without additional approvals.
- Agile and DevOps environments tend to “forget” about change management due to the speed of deployment expected by the consumer. In fact, the way in which Agile handles user stories is simply change management as long as it includes assessment of the requested functionality, impact on existing services and reduction of conflicts with other user stories. In DevOps, where as much as possible is automated, these same steps can be performed, possible in an automated way as well, depending on the nature of the services. Most of DevOps automation is in the area of release and deployment management, though, which will be covered next.

Release and Deployment Management

Release management and deployment in Service Management are closely related but distinct processes.

Simply put, release management is about what and when to release a service or service element, focusing on the bigger picture and ensuring alignment with business objectives, whereas deployment is about the how, dealing with the technical specifics of getting a new or updated service live and operational.

Objectives

The main objectives of release and deployment management are to plan, schedule, and control the movement of releases to test and live environments. Its primary goal is to

ensure that the integrity of the live environment is protected and that the correct components are released.

Implementing a well-structured release and deployment management processes helps minimize disruptions, improving service quality, and ensuring that changes are delivered in a controlled manner.

Benefits

The main benefits of good release and deployment processes are as follows:

- **Improved service quality**: by managing releases effectively, the negative impact of changes on the service quality can be reduced;
- **Improved risk management**: identifying and managing risks associated with releases helps minimize disruptions to the live environment;
- **Increased efficiency**: streamlining the release and deployment processes leads to more efficient use of resources.

Key Concepts

Release management includes the activities focused on the planning, scheduling, and controlling of the movement of releases to test and live environments. Its primary goal is ensuring that the integrity of the live environment is protected and that the correct components are released. It encompasses the oversight of multiple changes and features being bundled together into a single release, ensuring that they are integrated and tested in a pre-production environment before being deployed. This process, often documented in a release plan, involves coordination between different teams, planning release windows, documenting releases (e.g. release notes), and communication and training to stakeholders.

Deployment, on the other hand, is the process of moving the release into a different (staging, test or live/production) environment, or into the hands of users. It's more technical and involves the activities necessary to implement the new or updated service into operation. This includes preparing the target environment, installing components, configuring them to match live settings, and performing any necessary testing to ensure that the deployment meets required standards.

A few other concepts in these processes are:

- **Release Policy**: A set of rules that guide how releases will be managed within the organization, including release timing, training requirements, documentation needs, communication about releases, etc.
- **Release Package**: A set of configuration items (CIs), hardware, software, documentation, etc., that are to be released together.

- **Deployment Plan**: A detailed plan that outlines how the release will be moved into the live environment.
- **Environment**: Refers to the live, test, and development partitions or areas where services are built, tested, deployed and managed.

Process

1. **Plan the release**: define the scope, schedule, and resources required for the release.
2. **Build and test**: develop a release package and conduct thorough testing to ensure it meets functional and non-functional requirements. Consider testing user, security, quality and other requirements.
3. **Release readiness check**: ensure all aspects of the release are complete and ready for deployment; this may also include ensuring user documentation is updated, and that support teams have been trained, or are otherwise ready to support the release in the target environment.
4. **Deploy**: move the release package into the live environment according to the deployment plan.
5. **Review and close**: after deployment, review the release process to identify any lessons learned and formally close the release. This may include the closure of any related problem tickets, or incidents the deployment is meant to resolve. Any changes bundled into the release should now also be updated.

Common Pitfalls

- **Lack of communication**: poor communication between teams can lead to misunderstandings and errors during the release process;
- **Inadequate testing**: insufficient testing can result in releases that cause more issues than they solve;
- **Poor planning**: without thorough planning, releases can suffer from delays, resource shortages, and scope creep;
- **Failure to manage risks**: not identifying and managing risks can lead to unexpected issues during deployment;
- **Inadequate documentation**: failure to properly document the release and its components can lead to confusion and difficulties in troubleshooting.

Provide

The third stage in the ITSM.express approach includes a comprehensive approach to protecting, measuring, and enhancing IT services. It is here that the services that were designed, built and delivered are to be protected from threats in the operational environment. Many of these threats would be identified and logged in a risk register in the Define Stage, although new threats can and should be identified and managed in all stages. In this provide stage, it should also be ensured that the service provider is living up to promises made in earlier stages, by measuring service and service provider performance. In this stage we are also called to commit to continually improve the services offered to our consumers; improvement opportunities can be identified through the measurement of the services, and also through innovation.

Protect - Information Security

Objectives

When it comes to information security, a risk-based approach should be taken using the guidance in risk management in the Define stage. Again, here the ISO 31000 standard is used as high-level guidance. An important note regarding information security is that everyone in the service provider organisation, suppliers, partners and consumers should contribute positively to information security. The information threat landscape is constantly evolving, so it is important to stay ahead by regularly updating risk assessments and mitigation strategies. This proactive approach helps maintain an effective information security posture and also ensures compliance with relevant information security and privacy regulations.

Note that ISO/IEC 27001 is the international standard for information security. It has many more requirements and details on running an information security management system than what is presented here.

Process

In order to identify information security risks, it must first be understood what it is that should be protected. It is important to identify and record the information security assets, with an understanding of how critical or valuable these assets are to the service provider and to the consumers. Information security assets can include physical information (e.g. paper records), as well as digital and personal information. This allows the service provider to apply appropriate levels of protection to information security assets.

Once it is known what information security assets need protection, a thorough risk assessment should be conducted, identifying threats that these assets might face. It is important to note that threats are not always malicious: information assets are also

vulnerable to acts of nature (fires, floods, humidity) and accidental threats. Per the risk management approach described, evaluate the potential threats and vulnerabilities that could impact these assets. By understanding the likelihood and potential impact of different threats, any necessary treatment of information security risks can be prioritized accordingly.

Treatment involves developing a strategy to mitigate these risks by implementing robust security controls. This would include establishing clear policies and procedures for managing information securely, implementing logical and physical access controls, and other controls like malware protection and backups. Information security controls would also include training employees on standard operating procedures (to build competence and prevent mistakes) as well as training on security best practices.

An important sub-process in information security is access control, which is both part of physical security (building access, secure areas) and application security (access to applications and servers).

Access Control

Access control refers to the processes and technologies used to manage and monitor access to network resources, applications, and data. It is used to ensure that only authorised users and systems can access certain information and functionalities, enhancing security and compliance. Access control is critical for maintaining the confidentiality, integrity, and availability of information. By controlling who has access to what, organizations can protect sensitive data from unauthorized access, prevent data breaches, and comply with regulatory requirements.

Implementing access control involves defining policies, roles, and permissions that dictate what resources users can access and what actions they can perform. This typically includes user authentication, authorization, and audit processes, often supported by identity and access management (IAM) systems, or related technologies.

Measurement

Objectives

Measuring IT services effectively is crucial for understanding performance, improving service delivery, and aligning with consumer needs and business goals.

In the Define section consumer reporting was described and the service provider should also measure the services and their performance to confirm and report that they are (or not) delivering the required service levels.

Benefits

The insights that are gained from data analysis should assist the service provider in making informed decisions. This could involve addressing areas where performance is lagging, reallocating resources, or making changes to services or processes. This indicates the contribution that measurement makes to the next process described in Provide, continual improvement.

Key Concepts

Services are measured for many reasons, including improving consumer satisfaction, reducing downtime, or optimising resource usage. In Define consumer reporting was described and the need to establish clear objectives, specific measurable metrics and their targets. Measurement is critical to reporting. There are many technical services, not seen by the consumer, that contribute to customer-facing services that should also be measured. These measures can contribute to performance insights and to identifying significant events that would require intervention to restore normal performance and service levels.

Process

As part of this process it should be determined what data to collect, and how this data will be interpreted and used for the performance metrics, including choosing the right tools and processes. For example, the service provider can use automated monitoring tools to track system performance or use surveys to measure consumer satisfaction. Measuring activities, especially when automated, can generate large volumes of data so it is important that it is understood what data is being collected and why it is: what is the value of the data for the organization's objectives?.

Once the necessary tools are deployed, technologies and surveys to collect data, the service provider should ensure that these tools are integrated properly with dependent systems and that they are capturing data accurately and consistently.

The collected data should be regularly analysed to gain insights. You should look for trends, patterns, and areas of concern. This step often involves using data analytics tools that can handle large volumes of data and provide meaningful visualisations, and may contribute to, or even form the basis of event management activities and tools.

Summarised findings highlighting concerns, exceptions and trends should be included in clear and concise reports. These reports should be tailored to their audience – for example, technical teams might need detailed performance data, while executive management, service owners and managers, and consumers might prefer high-level summaries.

Common Pitfalls

Review and Adjust: IT environments and business needs are constantly changing, so it's important to regularly review and where appropriate adjust our metrics, collection methods, or analysis techniques. This ensures that our measurements remain relevant and aligned with our business objectives.

Improve

Objectives

Like risk management and information security, everyone in the service provider organisation, suppliers, partners and consumers can contribute positively to continual improvement. The organization should learn from customer feedback, and the data collected, analysed and reported on in Measure, to assist in making iterative changes to enhance service delivery, better meet service targets and business needs, or improve the service targets themselves.

Process

A common and straightforward approach to continual improvement involves the following steps:

1) **Identify the improvement opportunity**: Recognise areas where our service isn't meeting expectations, where processes or operating procedures are inefficient, or where other opportunities to improve exist. These opportunities could be identified through customer feedback, performance metrics, risk management, information security management or regular service reviews.
2) **Understand the current state**: Examine the existing service, process, or other target for improvement to gain a proper baseline prior to improvement, and to understand what improvements can be made. This could involve examining incident reports, workflows, or resource allocation.
3) **Set clear and realistic objectives**: Once it is known what needs improvement, the organization should set clear, measurable goals that are achievable. Small frequent iterative improvements often show great results with minimal disruption or resistance. Be clear about what you want to achieve with the improvement. This could be faster response times, fewer errors, or higher consumer satisfaction. At this step it should be determined how the success of our improvement will be measured and expected targets should be established. It makes sense to use the same metrics or feedback methods that identified the need for improvement in the first place.
4) **Plan the improvement**: Based on the analysis, develop a plan for how to make the improvements. This should include specific actions, resources required, and a timeline. It is important to be realistic and to consider potential challenges.

5) **Implement the improvement**: Put the plan into action. This may involve training staff on new procedures, updating hardware or software, tweaking or removing process steps, adding controls, or changing communication methods with customers.
6) **Monitor and review the improvement**: After implementing changes, closely monitor the service to see if the improvements are working, measure the improvement using the metrics and targets planned when we set the objectives for the improvement. Based on monitoring, it may be needed to make further adjustments to the improvements made.

Service improvement is not a one-time thing, it is ongoing, it is a continuous cycle of improvement, and so the next step is to pick up with the first step again, identifying the next improvement target.

Respond

Service Providers frequently have to communicate with consumers as part of service provisioning—specifically the end users of the services, irrespective of whether the Service Provider or the user initiates this communication.

It is important that there are clearly defined means of communication, which also provide an audit trail. The best way to do so is to provide a single point of contact, such as a service desk; however, this should not be misconstrued as a single *means* of contact.

Respond processes are directed by the communication facilitated by the single point of contact, and the systems put in place to facilitate this communication provide a means of escalation, activity and status tracking, and day-to-day measurement and management.

Traditional processes found in the RESPOND stage are:

1. *Request Management*
2. *Incident Management*
3. *Problem Management*
4. *Service Reporting, and*
5. *User Update/Alert Management*

The Response process outlined in what follows presented here as a unified approach to dealing with user queries, including Incident Management, Request Management, and Problem Management. We call this unified process the Response process. The single point of contact for communication around the response process is called the Service Desk.

The Service Desk

Objectives

The objectives of a Service Desk are to:

- Provide a single point of contact for service consumers (primarily users);
- Facilitate service restoration;
- Direct user requests and facilitate outcomes;
- Provide first-level resolution of incidents and requests where possible;
- Log and track communication and intervene if process activities do not meet quality standards (typically specified in an SLA); and
- Be a channel to effectively handle issues when the provider depends on the support of third parties (e.g. escalation management).

Benefits

- Users know exactly who to contact and how to contact the service provider;
- Issues (primarily incidents) and service requests can be recorded and therefore managed to ensure service standards are adhered to;
- Data recorded can be used to:
    - improve services and service components;
    - report on service provider performance; and
    - better understand the state of service to improve future service value.

Process

A Service Desk typically:

- Records user queries (Incidents and Service Requests);
- Responds to user queries by providing first-level support;
- Communicates with third parties the service is dependent on;
- Provides users with feedback;
- Keeps stakeholders updated (e.g. during service outages); and
- Reports on data collected to key stakeholders.

The second bullet is of particular importance, as a service desk agent should have the technical skills and tools to resolve and close the majority of queries (incidents and requests) logged at the desk.

The Service Desk (especially in small organizations) may also be tasked to perform other operational activities like service monitoring, doing backups, and managing operational support activities.

Common Pitfalls

some common issues occurring when operating a service desk and how to fix them are the following:

1. A lack of business (user context) understanding

Service desk staff must understand how the consumer's business operates and immediately identify if a user query relates to a critical element of business success.

Service desk Agents need to receive basic training to help them understand the contexts of queries they may have to deal with during their duties.

2. A lack of technical understanding

Service desk agents must understand the design of the services and the dependency between service components and the associated technical skills to be able to provide first-call resolution to the majority of issues logged at the service desk.

Service desk Agents must have a high-level view of the architecture of their services. This helps them to quickly identify the next points of escalation if they themselves cannot resolve an issue.

Ensure that service desk agents have sufficient knowledge to deal with most user queries at the first point of contact. Train service desk agents on the technical environment they need to support and provide them with the right tool to resolve user queries.

3. A lack of correct information/data and integrated tools

A service desk cannot function without proper support tools, and because service desk tools are often used for support and technical management functions, careful consideration, selection and investment are required. An integrated service management system is required to do this work.

Support will always be sub-optimal unless data recorded about user queries can be effectively managed across the provider landscape. Three elements are of particular importance:

- The correct data needs to be recorded to facilitate service quality outcomes;
- Workflow must be simple and configurable to align with the type of query being dealt with;
- Data about all recorded queries and access to other data that will assist in the resolution of a query must be freely available and easily accessible.

4. A lack of communication skills and empathy

Service providers should invest in training that helps agents communicate better and empathize with users. Effectively facilitating outcomes in this scenario is not only about technical ability; user satisfaction often exclusively depends on their interactions with the service desk.

Measuring successful outcomes

Measuring service desk performance can be done in several ways and is closely related to the measurement process and to continual improvement in the Provide phase.

Possible measures for the service desk from a user perspective are:

- Ease of contacting the Service Desk;
- Quality of query resolution;
- Speed of query resolution;
- Experience with Service Desk interaction.

Possible measures for the service desk from a service provider perspective are:

- Time to respond to a query;

- Time to resolve the query;
- % First call resolution achieved;
- Number of calls per (time period) per agent handled;
- Number of calls per (time period) per agent resolved;
- % of calls re-opened – although some would disagree with the practice of re-opening calls, it is important to ensure that once calls are closed the query was successfully resolved, and re-occurrence is minimized – this is what is measured here.

Resolving User Queries

A simple single process is presented here to handle all user queries - which may, in traditional service management terms, be incidents, service requests, and even problems.

What is represented here is a simple baseline that can always be expanded on as organizational needs evolve; however, complexity brings a lot of additional challenges.

User queries are often not fully dealt with, as workarounds or temporary fixes are put in place to rapidly restore the service, knowing that the root cause of the outage was not addressed. The guidance in this section is based on the experience that most organizations pay lip service to addressing root causes after the fact, and as such a 'single-process' approach is outlined here.

The Response Process

Objectives

The objective of this process is continued availability of services. The process ensures that by facilitating communication and user requests, restoring service outages, and ensuring service outages do not re-occur.

Benefits

- Improved productivity of service end users;
- Improved productivity of service provider staff;
- Improved satisfaction of service end users;
- Accurate data on service degradation or outages and facilitating service design and planning;
- Facilitating a better understanding of the cost of service unavailability and the true cost of the service.

Key Concepts

This process deals with three key concepts and lines of work, and providers who think that they want to separate this into three distinct processes should feel free to do so.

However, we believe that there are so few differences between how incidents and requests are dealt with that in many organizations, the distinction is academic.

Over decades, we observed that when incident (e.g. service unavailability or degradation) handling becomes the primary focus, dealing with the root cause after service restoration (generally referred to as problem management) is neglected. Therefore, we believe that an integrated approach is justified and, where it has been implemented, is prudent.

But what are the incidents, problems, and requests that this process deals with?

Incidents are service disruptions or the degradation of a service that prevents a user from obtaining the full benefit that the service is supposed to offer. These instances of disruptions or reduced service performance are often the result of some fundamental structural issues within the design or operation of the service, which are often not obviously visible. Still, unless these issues are addressed properly, they will reoccur because the root causes of the incidents remain. **Dealing with the root cause of incidents is often called problem management.**

Any other query a user has regarding using the service is a request. This includes questions about its use, access to the service of service features, or service provider actions to help users maximize the benefits of using the service. These often include requests for customized service output, consumables used in service provisioning, or special output - like service reports.

Note that requests to provide new or changed service features are not handled as service requests but rather as requests for change and are facilitated by the change management process. This process may be used to record such requests and then pass them on to the process dealing with service change and customization.

However, many of the changes users request are often well understood, and the risks are known. They can be managed by fulfilling the request within certain boundaries, and frontline support staff can be taught how to facilitate that standardized change.

Organizations assign **prioritization levels** to requests based on the **impact and urgency** of the request and the associated timelines that influence the realization of the negative effect**.**

Services are often restored using a temporary fix or workaround and, conceptually, an incident is closed once the user can resume their normal daily functions.

It should be noted that even if it seems that the incident is closed, the root cause has not been dealt with and this often results in more service degradation, unstable service, and more service disruptions – which soon may become a vicious circle.

In a traditional service management approach, a separate process (problem management) is used to find root causes and implement permanent fixes – the only problem is that *very* few organizations actually do this well or at all.

We propose an integrated incident/problem management approach, and evidence suggests that the likelihood of the causes of incidents being dealt with is much higher using an integrated approach.

All common approaches to finding and eradicating the root cause of one or more incidents have the following in common.

1. Defining what the issue is;
2. Finding the root cause - note that causal analysis is not necessary; root-cause analysis, as a cause, can have a cause itself: to fix an issue permanently, you need to trace the cause of causes until you are at the root of the problem;
3. Implementing measures to fix the root cause.

A note of caution is appropriate here – do not assume that a single root cause exists – there are often root cause**s.**

Not all user requests should include root cause assessment, but those with a high priority definitely should.

Process

The overall process to ensure a comprehensive response to service queries involves the following components.

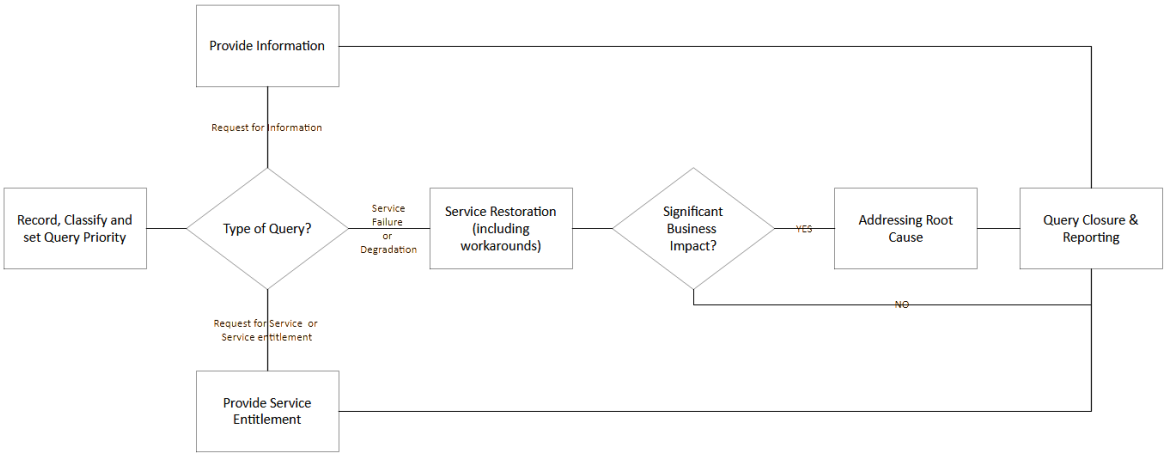

1. **Record, classify, and set request priority**

When user requests are logged, service desk agents should understand the services well enough not only to record the query but also correctly to classify the type of request and assign a priority to the request that is aligned with business impact or effect.

Recording usually involves information like user details, location, affected service, and typical symptoms or effects experienced.

Classification involves deciding what type of request it is (for example, a request for information, request for a service entitlement, or service degradation or failure – used to trigger a specific workflow), and lastly, the effect on the business, which is used to define a request priority.

Based on the request type, a specific workflow or procedure may be used to enable speedy closure of the request.

2. **The Service Desk**

Just because the service desk is a single point of contact, it does not mean that the phone is the only source of access to the service desk. The service desk should be able to effectively facilitate user requests regardless of the means of contact the user made.

3. **Provide Information**

Organizations often prepare a standard response to commonly asked questions (FAQ), a handy source for answering user requests. This can be made online and should also be available to all service desk agents.

If a user's request for information cannot be satisfied by providing a standard answer. The service desk should refer the query to the resource in the organization that would be best placed to do so - once again, business knowledge is needed to do this effectively.

4. **Service Entitlements**

Service entitlements are permissions for end users to access resources, such as application access or access to a document. This query type can easily be handled at the service desk once the agent validates the entitlement.

It is, however, sometimes also necessary to escalate service entitlements to an individual with the skills, access, and seniority to deal with this type of query specifically.

5. **Service Restoration**

Restoring failed or degraded services often requires technical competence; however, deep technical knowledge is not always required if previous efforts are properly recorded and used as a reference database.

Define the possible solution to commonly occurring service issues at the service desk in an attempt to improve first-call resolution rates.

If the service desk agent cannot resolve the incident, they should have enough technical knowledge to forward (escalate) the query to the technical resource most likely able to resolve the incident.

If incidents are escalated, the service desk must keep tabs on progress and validate with users that the incident was dealt with to their satisfaction when the technical resource and the end user indicate the incident was resolved.

Here, a sub-process is inserted, usually defined as a separate process called problem management.

Before incidents can be closed the service desk should determine whether the incident is of a critical nature with a significant business impact or not. If the answer is affirmative, the incident cannot be closed but should be assigned to a more senior technical resource for root cause analysis. In this instance, the incident can only be

closed once the root cause of the service degradation or failure is identified and addressed.

6. **Addressing Root Causes**

All service degradation and failures that had a material effect on the consumer's business must be further investigated (after service restoration) to determine the root cause of these outages. Since these outages had a material effect on the business, the service provider cannot neglect to ensure that the service outage or degradation does not re-occur in the future.

The method chosen to do so is up to the service provider.

The only specific guidance that remains here is that this task is unlikely to be assigned to an individual, as no one person would have the requisite skills to look at the problem from different angles. Problem solving is best done using a multidisciplinary team of technical and business experts to truly get to the bottom of the root cause or causes of the service degradation or outage.

Once root cause(-s) are found, actions should be taken to address the issue permanently. This often involves actions that require a formal assessment of making these changes and should be handled by the Change process.

7. **Closure and Reporting**

All requests should be formally reviewed and closed; this can often be as simple as asking the user who logged the request if they are happy with the response and resolution offered.

However, a formal review and the implementation of corrective action (addressing the cause(-s), usually subject to change management) are the final closure criteria for service degradation and failures that had a material effect on the business.

In these cases, it is also prudent to evaluate query trends to validate that the area addressed has shown a marked improvement in terms of the quality of service provided. Utilizing trend analysis to evaluate service availability and capacity improvement.

Access to the service desk and the resolution process data should be accessible and anonymized where private information is concerned before being shared by other service disciplines.

More detailed guidance will be provided under the service reporting process .

Common Pitfalls

- Respond activities are only effective if a disciplined approach to logging, tracking, responding, and resolving issues is maintained in the organization. This

often means significant education of customer and service provider staff. Everyone should understand what is done, how it is done, and what to expect.

- Service desk agents who cannot resolve most queries at the first point of contact are not helpful and do not provide the correct level of service. The typical catch-and-dispatch approach adopted by many organizations should be avoided – this is not a Service Desk, as it does not have the ability to provide service value to users and other consumers. Service desk agents should be well-versed in this process and have the correct skills, information, and technology to facilitate smooth user query resolution.
- Methods of communication should suit the context in which they are deployed. Be careful when deploying self-help logging of incidents where user skills and the understanding of the service context are not significant. Sending an email where someone from the service desk needs to phone back and re-log the query is not productive and prolongs positive service disruptions.
- Ensure that the root causes of all service degradation and failures that had a material effect on the business are investigated further and a permanent fix is put in place to address the cause(-s).