SOP 0020 - Incident Management

Purpose: The Faculty of Medicine Digital Solutions follows the UBC IT Incident Management process as outlined below which was written by Joanne Wester from the ITSM (IT Service Management) department.

Introduction

The Value of a Common Process

UBC is a diverse environment, and IT has evolved in numerous faculties and departments to meet the needs of local end users, resulting in myriad processes for handling user and system needs. Despite this apparent diversity, it is clear that:

- At UBC, regardless of the faculty or department affiliation, customers have similar types of technology service needs (they want something fixed, they need access to a system, they want instruction on how to do something, etc.)
- At UBC, the majority of customers are served by multiple IT service providers; IT groups need to co-ordinate activities in order to deliver services to customers
- At UBC, IT groups have a shared interest in supporting their customers' needs with technology services
- At UBC, IT groups conduct similar types of work in order to support their customers

Given these similarities, documenting and sharing standard and repeatable processes will enable IT groups across UBC to mature our ability to serve our customers. This document outlines processes based on ITIL v3, but customized to work at UBC by members of various faculties and departments. The ultimate goal is to provide individual IT groups at UBC, regardless of faculty, size, or budget, the agility to serve their local user populations better.

What is Incident Management?

The goal of Incident Management is to restore normal operations as quickly as possible with the least possible impact on business processes and/or end users. Inputs to Incident Management mostly come in the form of Incidents raised by end users, but may also be the outcome of event management and operational support activities. Incident Management also refers to managing the lifecycle of all Incidents. An Incident lifecycle is the end-to-end process from when the incident is reported or suspected, until the incident is resolved or a workaround is put in place.

The Incident Management process described in this document is based on the Information Technology Infrastructure Library (ITIL) v3 framework, and is technologically independent i.e., the process activities described can be implemented in any technological environment, and are applicable to any organization following an Incident Management framework. Furthermore, the descriptions contained in this document are to serve only as a blueprint or guidelines, and are not meant as a "how to" or as a manual.

This document is based on the larger UBC Incident Management Framework, and provides a high level view of the Incident Management process and serves to act as a Quick Reference Guide for various IT groups at UBC.

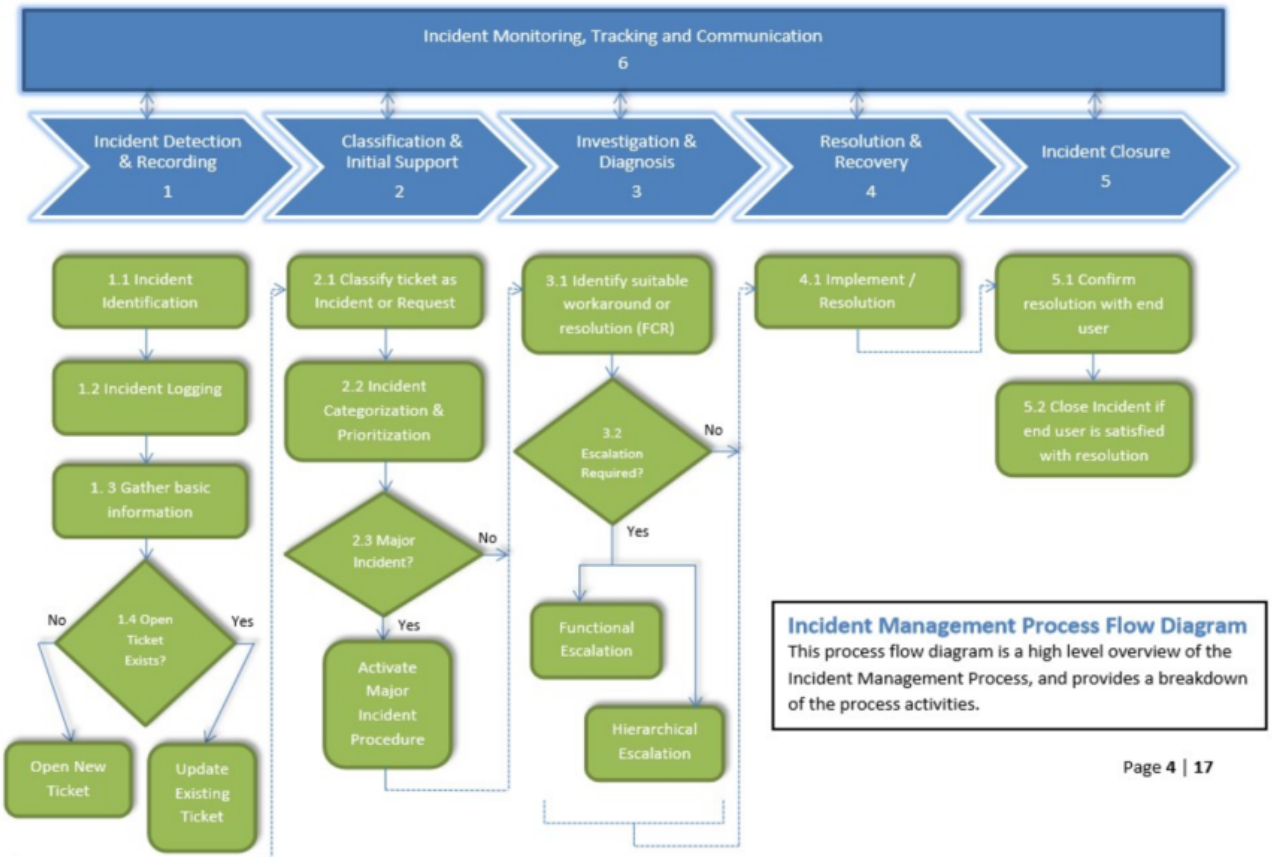

Incident Management Process Flow Diagram
This process flow diagram is a high level overview of the Incident Management Process, and provides a breakdown of the process activities.

Guiding Principles

In order to improve process efficacy, the Incident Management Working Group and Community of Practice agreed to a set of principles and supporting best practices outlined below to inform the incident management process across UBC.

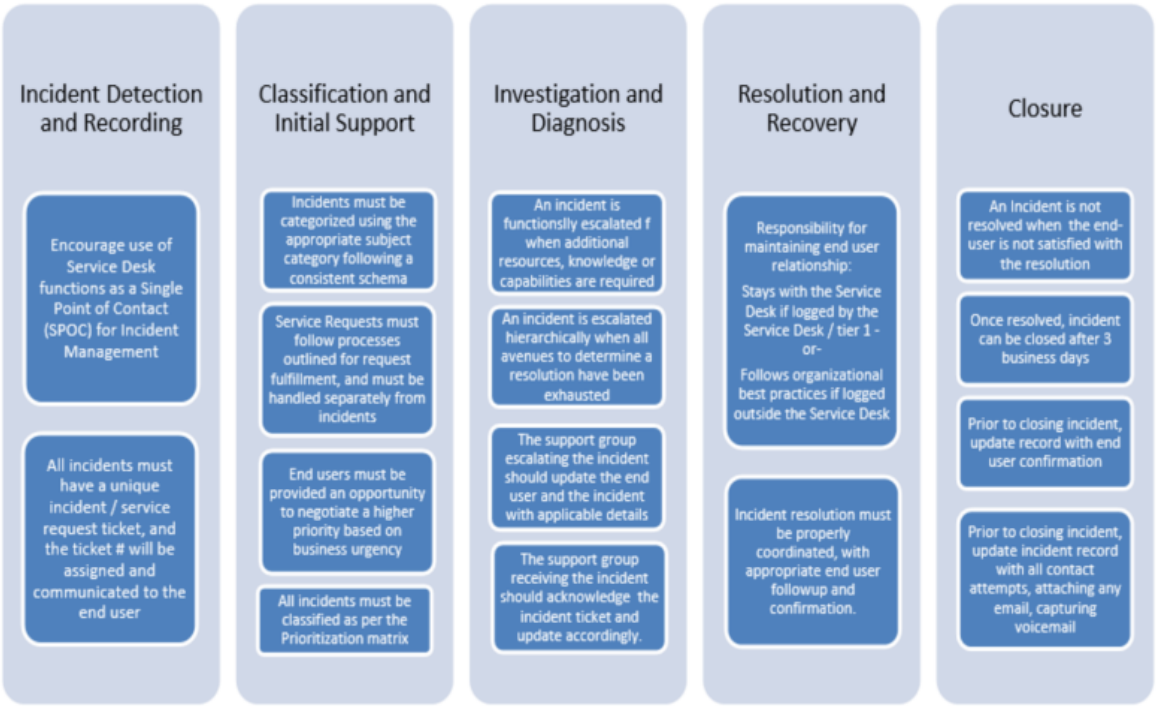

Incident Management Process Activities

1. Incident Detection & Recording

1.1 Incident Identification:

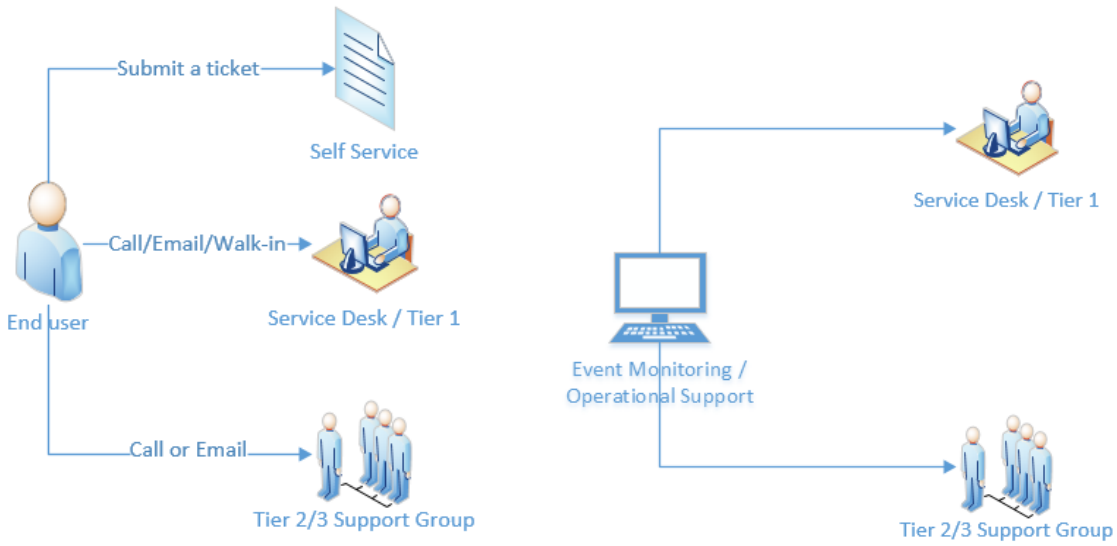

Incidents are identified in one of two ways:

1. Reported to tier n by end users
    a. End users submitting a ticket via Self Service
    b. End users contacting tier 1
    c. End users contacting tier 2/3
2. Suspected by tier n based on operational support / event monitoring activities

1.2 Incident Logging:
Once an incident has been reported, support analysts must immediately log the incident to mitigate any risk of new incidents emerging which may cause the process to be postponed indefinitely. Service Desk / tier 1 must establish whether the service(s) required are supported and included in the end users' Service Level Agreement (if applicable), and if not:

1. Transfer end user to the appropriate Service Desk via warm transfer, or,
2. Provide information to the end user regarding supported services, and where / how to obtain further support

1.3 Gather Basic Information:
Support analysts must gather and record basic information related to the incident such as time, description of the incident, systems affected etc. based on reports from the end user, or based on event management system/software. The information collected also includes the contact information for the end user/client and the organizational group, department, faculty etc.
1.4 Open Ticket Exists?
The incident must be assigned a reference number to uniquely identify it in the ITSM system, and when communicating with end users. To avoid duplication and to ensure that incidents that have been raised previously are accurately updated, support analysts must determine whether or not an open ticket already exists for the reported / suspected incident and either update the open ticket or create a new ticket including the information gathered in Step 1.3

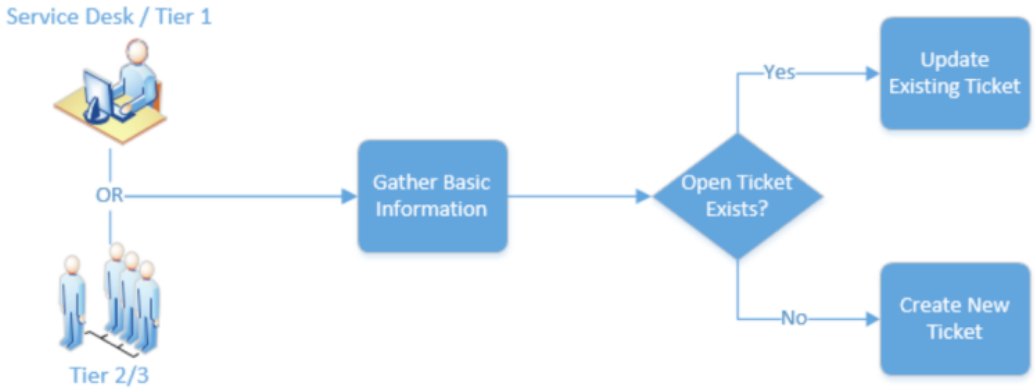

2. Classification and Initial Support
2.1 Classify ticket as Incident or Request:
Based on the information provided by end users, support analysts classify a ticket as an Incident or a Request:

- - **Incident**: An unplanned disruption in the standard delivery of an IT service or operation, or an event that may cause interruption or degradation in the quality of a service. Incidents follows a break / fix cycle, e.g., issues with network connectivity to a building / facility
- - **Service Request**: Repeatable tasks that are frequently requested by end users, e.g., provisioning access to a system

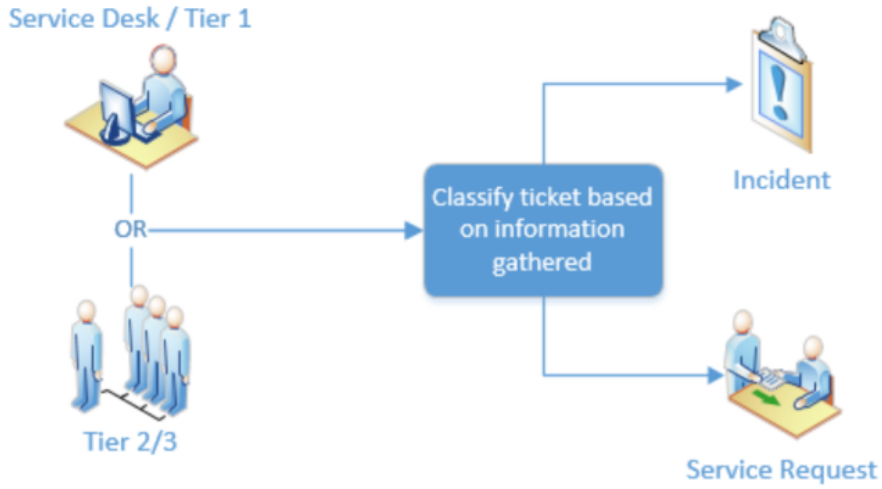

2.2 Incident Categorization and Prioritization:
Support analysts must categorize incidents by selecting the appropriate Category from the subject tree (subject/type/item).
Incidents are prioritized based on the assessment of Impact and Urgency.

- - **Impact**: A measure of the extent or effect of an Incident on business processes. Impact is often based on how Service Levels will be affected
- - **Urgency**: A measure of how long it will take to fix / resolve the issue and has a significant impact on the business.

Once the Impact and Urgency of the Incident are selected, a relevant Incident Priority code is assigned to the incident based on the Incident Prioritization Matrix (see Appendix).
2.3 Major Incident
Once the priority of an incident has been calculated, the Service Desk / tier 1 must determine whether the incident qualifies as a Major Incident, based on which the Major Incident Procedure is activated. For critical incident, please refer to the FOM DS major incident process.
3. Investigation and Diagnosis
3.1 Identify suitable workaround or resolution
Once Incidents are logged, categorized, and prioritized, support analysts troubleshoot the underlying issue following basic troubleshooting techniques, and work towards identifying a suitable resolution, or establish a workaround that allows resumption of normal business operation.
If tier 1 support analysts are able to resolve the incident without requiring any support from tier 2/3 groups, the incident is said to have been resolved on first contact. The established metric for this is referred to as First Call Resolution (FCR).
Support analysts must confirm incident resolution with end users before they record the incident as having been resolved. If end users are unattainable, support analysts must try to establish contact (via phone or email) as per organizational support policies and procedures. If the end users cannot be reached after repeated contact attempts, the incident can be closed.

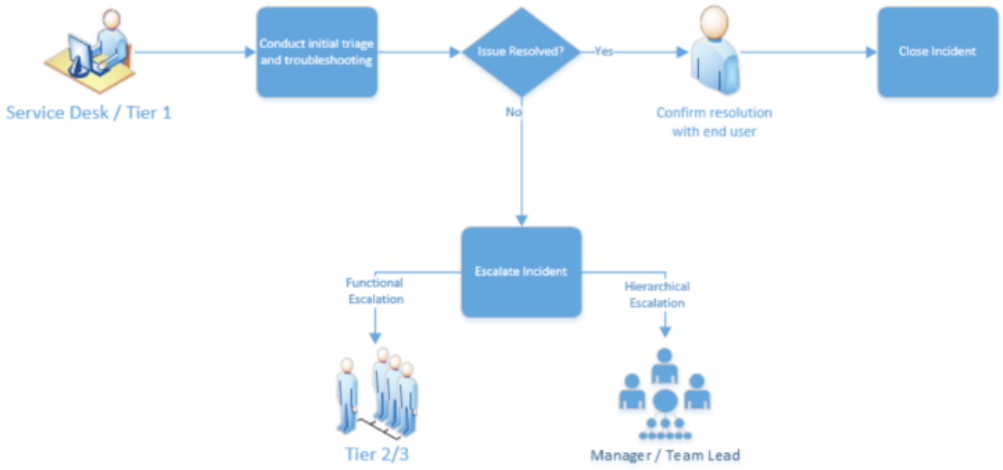

3.2 Escalation Required?
If tier 1 support analysts are unable to resolve the incident or establish a workaround on first contact (FCR), they can escalate the incident to the relevant tier 2/3 support group. There can be two types of escalations:

1. **Functional Escalation**: If tier 1 lacks the domain-specific or technology-specific knowledge to resolve an incident, the incident can be escalated to a tier 2/3 group for troubleshooting / resolution.

2. **Hierarchical Escalation**: If managerial approval is required to resolve an incident, tier 1 can escalate it to the appropriate organizational level to action the incident.
3. Off-hours emergency escalation: Please create a ticket in ServiceNow and call the Operations Center at 604.822.6141 as they are open 24/7 and can engage with the service owner on your behalf. If you are unable to reach the Operations Center, please escalate to Jayson To, Service Delivery Manager, at 604.889.5023. Note: The Operations Center is different from the ITSC helpdesk whose operating hours are from Monday to Friday, 07:30 to 18:00 Pacific Time.

4. Resolution & Recovery

4.1 Implement Resolution

Once a tier 2/3 team has identified the resolution to an incident, or a workaround has been established, tier 1 can facilitate the implementation of the proposed resolution.
Once a resolution has been identified and implemented, support analysts must contact end users to determine if the issue has been resolved, or the impacted services have recovered. This is to ensure unambiguous resolution of incidents for both IT and the end users.

5. Incident Closure

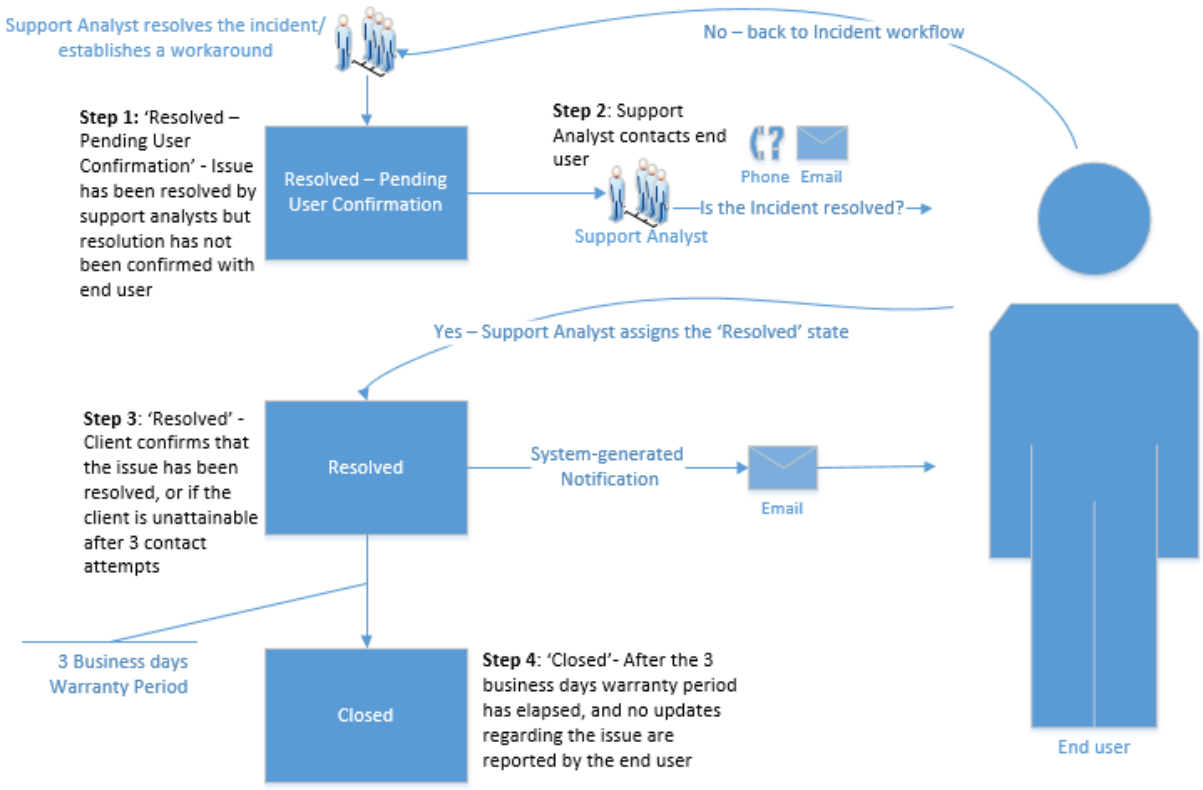

5.1 Confirm resolution with end user

Once a resolution has been identified and implemented, support analysts must contact the client to confirm resolution of the incident before it can be recorded as having been resolved. If end users are unattainable, support analysts must try to establish contact (via phone or email) as per organizational support policies and procedures. If the end users cannot be reached after repeated contact attempts, the incident can be closed.
Based on support procedures, tier 2/3 can reassign tickets to tier 1 in order to conduct the end user follow up and confirm resolution of incidents.

5.2 Close Incident

If end user is satisfied with resolution: Once a resolution has been put in place, and has been confirmed with the end user, the incident can be marked as Closed.
After being marked as Closed, the incident will remain in this state for a warranty period of 3 business day. The end user can request the incident to be re-opened within a 3 day warranty period. This is to ensure that end users have sufficient time to notify IT if the issue recurs or persists, or if they feel the incident is not fully resolved.

6. Incident Monitoring, Tracking, and Communication

Continuous monitoring and tracking of the processing status of incidents is important in establishing if the incidents will be resolved, or appropriate workaround put in place within an expected time frame as outlined in the Service Level Agreement (SLA). This ensures that counter measures can be introduced if service levels are likely to be breached.
Incident tracking also enables reporting on past incidents including Key Performance Indicators (KPIs) and metrics related to Incident Management. This ensures that support groups have sufficient knowledge of any important trends, and are prepared to respond to any anticipated future increases in incident volumes.
End users must be pro-actively informed of service failures as soon as they become known to the Service Desk / tier 1, so that end users have sufficient time to adjust in case such failures affect them. Proactive communication also helps to reduce the number of enquiries and reports by end users, and allows the Service Desk / tier 1 to focus on mitigating risks due to service failure / outage, instead of handling queue volumes.

Incident Management Critical Success Factors and KPI

Incident Management Critical Success Factors (CSF) include:

- Resolve Incidents Quickly
- Ensure End User Satisfaction
- Maintain IT Service Quality
- Improve IT and Business Productivity

The following matrix illustrates some of the mechanisms for measuring performance. By aligning to the listed Critical Success Factors (CSF),

benchmarking can be done for the associated key Performance Indicators (KPI).

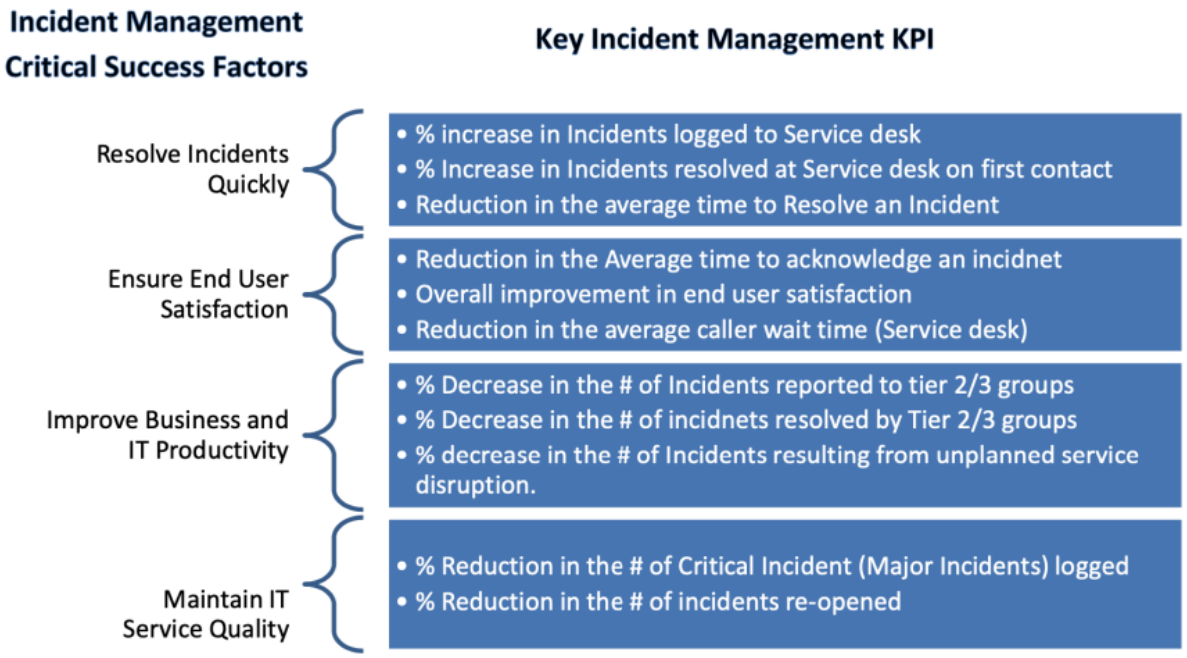

Appendix

Incident Prioritization Matrix

Prioritization occurs as a result of assigning a level of Impact and Urgency to the event. In the event of a service disruption an incident is prioritized by assessing the disruption's impact on the organization and urgency for resolution.

Assessing Impact

Impact is a measure of the effect of an Incident, Problem or Change on Business Processes. Impact is often based on how Service Levels will be affected. (Example # of user)

IMPACT – Scope of incident on organization		
HIGH	**MED**	**LOW**
'Campus Wide Outage' 'Significant # of customers impacted' 'Impact on individuals or services that will affect UBC's reputation, safety, or revenue'	'Individual department / faculty impacted'	'Single or sporadic customer impacted'

Assessing Urgency

Urgency is a measure of how long it will be until an Incident, Problem or Change has a significant Impact on the Business. (Example: Time of year)

URGENCY – Time Sensitivity and Business Criticality		
HIGH	**MED**	**LOW**
Business Criticality[1] = 'Mission Critical' = 'Core' Immediate response is critical	Business Criticality = 'Essential' Timely response essential	Business Criticality = 'Important' Immediate response is not required

Calculating Priority

Calculated PRIORITY				
		IMPACT		
		HIGH	MED	LOW
URGENCY	HIGH	Critical	High	Medium
	MED	High	Medium	Low
	LOW	Medium	Low	Low

Priority Based Service Level Agreements

Priority based Service Level Agreements (SLA's) measure the ability to acknowledge and resolve an incident based on its overall priority.
1. **Incident Acknowledgement Time** is *the time between when the ticket is created and the analyst's response back to the customer*, acknowledging receipt-(AnyState between = <Accepted> to <Closed> will stop the time for response time)
2. **Incident Resolution Time** is *the time between when the ticket is created and when it is "resolved"*
(State = Resolved or Closed will stop the timer for resolution time). Please refer to the FAQ for additional information on priority based SLA's

Priority	SLA - Incident Acknowledgement (*upon diagnosis and assessment)	SLA - Estimated Incident Resolution Time
1 (Critical)	Immediately	Immediately
2 (High)	<30 minutes	< 1 Business Day
3 (Medium)	4 Business Hours	< 2 Business Days
4 (Low)	1 Business Day	<5 Business Days

Unless otherwise stated in a SLA (Service Level Agreement) the *Incident Response Time* and *Estimated Time to resolve targets* are based on an organizations core business hours. If an incident occurs outside of core business hours, the expectation exists that it will be handled the following business day as per the targets. If there is a breach in the listed parameter the analyst will escalate according to support. For critical incident, please refer to the FOM DS major incident process (currently being finalized).

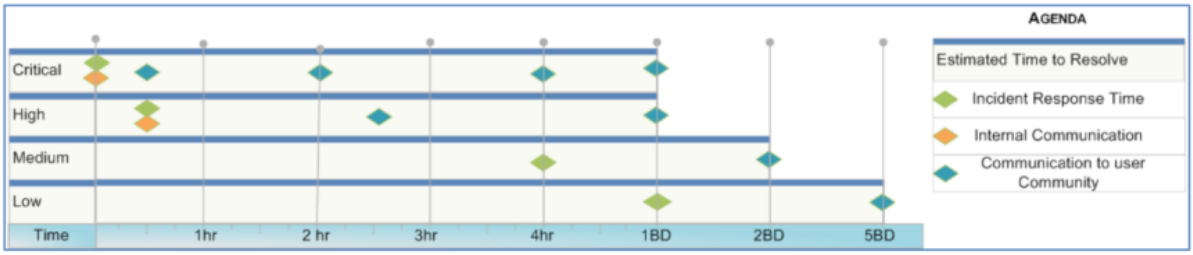

For each process activity in the incident lifecycle, a relevant status code must be assigned to an incident. This allows for efficient monitoring and tracking of an incident, and ensures clear and unambiguous communication of the incident status, both internally between support groups, and externally to the end user community. The table below lists the incident status codes for each process activity that correspond to 'close codes' in the Incident Management system currently in use by the organization:

Status Code	Description
NEW	The incident record should be assigned the '**NEW**' status while the incident is being identified and logged, and a new ticket is created (or an existing ticket is updated).
ASSIGNED	Once the incident has been recorded and has been classified and prioritized, the '**ASSIGNED**' status code should be used. This status should be used by the Service Desk / tier 1 during initial support.
ACCEPTED	The '**ACCEPTED**' status code should be used to indicate that initial information has been gathered, and the ticket is either with the Service Desk for First Call Resolution, or has been assigned to the relevant support group.
WORK IN PROGRESS (WIP)	Once a support group receives the incident record, the incident status should be changed to '**WORK IN PROGRESS**' to acknowledge that the identification of the root cause is underway, and efforts are being made to implement a resolution or workaround.
PENDING **CHANGE** **RECOVERY** **INFORMATION** **SCHEDULE** **VENDOR** **PARTS**	The **PENDING** status codes should be assigned if incident resolution is pending another task: **Pending Change –** On hold waiting for a change **Pending Recovery –** On hold waiting for supplementary recovery actions needed to completely resolve the incident **Pending Information –** On hold waiting for further information from the end user or a 3rd party **Pending Schedule –** On hold as further actions cannot be carried out until a scheduled time **Pending Vendor –** On hold waiting for information or recovery tasks from outside vendor(s) **Pending Parts –** On hold waiting for parts to complete a repair
RESOLVED – PENDING USER CONFIRMATION	When an incident resolution or workaround has been implemented, and support analysts believe that the incident is resolved, the **'RESOLVED – PENDING USER CONFIRMATION'** status code is used. This code also denotes that efforts to confirm the resolution with the end user are underway.
RESOLVED	Once the end user has verified and confirmed the resolution / workaround for an incident, the incident record is assigned the status code '**RESOLVED**'.
CLOSED	If no further updates are received from the end user within the 3 day warranty period, the incident record is assigned the '**CLOSED**' status. Once Closed, incidents cannot be reopened.

ITSM / ITIL Definitions

The purpose of this document is to provide a glossary of the terms utilized within the ITSM program.

Incident Management	The objective of Incident Management is to restore normal operations as quickly as possible with the least possible impact on either the business or the end-user, at a cost-effective price. The primary goal is to restore end-user functionality as soon as possible
Incident	An unplanned interruption to an IT Service or a reduction in the quality of an IT Service. Is any event which is not part of the standard operation of the service and which causes, or may cause, an interruption or a reduction of the quality of the service.
Problem Management	The process responsible for managing the lifecycle of all problems. The objective of Problem Management is to prevent Incidents from happening, and to minimize the impact of incidents that cannot be prevented. Problem Management aims to resolve the root causes of incidents and thus to minimize the adverse impact of incidents and problems on business that are caused by errors within the IT infrastructure, and to prevent recurrence of incidents related to these errors.
Problem	Is an unknown underlying cause of one or more incidents
Service Request	A request for information, or advice, or for a Standard Change or for Access to an IT Service. For example to reset a password, or to provide standard IT Services for a new user. Service Requests are usually handled by a Service Desk, and do not require an RFC to be submitted.
Event Management	The Process responsible for managing Events throughout their Lifecycle. Event Management is one of the main Activities of IT Operations
Service	• IT Service Offering should support a business process to an end user. It should include a service Level agreement outlining commitment around availability, monitoring and support function; • Must include professional services (e.g. evaluation, support, transition) and may be composed of collections of both HW/SW/ Data (Applications)
Customer	An individual or group negotiates with the service owner to fulfill goals of on behalf of the end-users of the service. Is responsible for user experience and service level targets.
Classification	The act of assigning a Category to something. Classification is used to ensure consistent management and reporting.
End-user(s)	• Executes a step or steps of the business process through interaction with the application or output from the application. • Experience the IT Service Offering as represented through the service catalog.
Priority	Sequence in which an Incident or Problem needs to be resolved. Priority is based on Impact and Urgency, and is used to identify required times for actions to be taken. (Impact +Urgency)
Impact	A measure of the effect of an Incident, Problem or Change on Business Processes. Impact is often based on how Service Levels will be affected. (Example # of end-user)
Urgency	A measure of how long it will be until an Incident, Problem or Change has a significant Impact on the Business. (Example: Time of year)

Service Level Agreement	Detail agreement that is negotiated between a customer and a service owner outlining the delivery, targets and expectations of the IT service. This may include details on items such as availability, support, cost, and contingency.
Service Level	Measured and reported achievement against one or more Service Level Targets. The term Service Level is sometimes used informally to mean Service Level Target.
Functional Escalation	A horizontal escalation where an analyst transfers an incident to another Tier when differing / more skills are required. ITIL defines functional escalation as "Transferring an Incident, Problem or Change to a technical team with a higher level of expertise to assist in an Escalation"
Hierarchical Escalation	Is considered a vertical escalation up the organization when more authority is required (example: senior management) ITIL defines vertical escalation as: "Informing or involving more senior levels of management to assist in an Escalation