

# Automate incident and change management

**Improve the availability and agility of enterprise services**

## What's in this Success Playbook

Automation can meet new expectations for incident and change management, and it delivers both the stability that the process needs as well as the speed that the new business drivers demand. But automation isn't simply push-button. This Success Playbook will help you:

- Build a framework that uses automation to identify and resolve critical service issues faster
- Build approaches that use automation to improve the accuracy, speed, and efficiency of change management

## Key takeaways

### The most important things to know

- Automation requires a solid set of service definitions and a framework the enterprise has agreed on that helps explain service criticality in terms of its business impact. Automating incident and change management is only helpful when you know where the greatest risks to your business outcomes are.
- Automation requires sufficient, credible data to support response, classification, and remediation.
- Automation should be iterative and focused on continuous process improvement. Few organizations are likely to realize the full potential of out-of-the-box automation and should exploit process performance data to identify where they can continue to fine-tune their approach.

## The payoff of getting this right

Effective automation can improve the speed, accuracy, and throughput of your incident and change management processes—helping you to manage cost and improve the quality of IT services.

## What you need to get started

### Prerequisites

You need a working knowledge of your current incident and change management processes.

## Playbook overview

To get the most from automated incident and change management using ServiceNow®, follow these stages:

- Stage 1** – Establish clear dependency mapping
- Stage 2** – Proactively identify service issues
- Stage 3** – Automate incident response and resolution
- Stage 4** – Automate change management
- Stage 5** – Measure impact and tune

## Stage 1 – Establish clear dependency mapping

### KEY INSIGHTS

- The benefit of automating depends on the quality and availability of the underlying incident and change management data required.
- Establish maturity in configuration management before you automate.

Before you begin an automation program for incident and change management, you first need to understand:

- How hardware and software assets are connected to deliver IT and business services
- How asset-level failures impact service delivery

To establish this foundational understanding, first ensure you have a mature approach to configuration management, supported by an effective discovery process.

### Configuration management

If you're using the ServiceNow Configuration Management Database (CMDB), you should have a single system of record for infrastructure, application, and service data. The effectiveness of your CMDB depends on how you approach planning and data model design.

ServiceNow recommends keeping CI classes and attributes simple and aligned to use cases and goals for configuration management. For example, if your goal is to gain accuracy in infrastructure configuration mapping so your team can understand change impact, align your CIs with that goal.

As a general rule, start using the out-of-the-box CI classes on the Now Platform™, which contain all of the attributes that are typically required to support incident and change management automation.

For more insight on effective configuration management practices, see our [Success Playbook on planning your successful CMDB deployment](#).

### Discovery

ServiceNow Discovery is an integrated and agentless means to populate CI data in the CMDB. Discovery automatically identifies the type of CI it interacts with and launches additional probes and sensors to gather further information and attributes. The CMDB automatically checks the

data for errors, normalizes and transforms it, and then loads it to ensure the most recent and accurate profile of that CI.

For more insight on effective discovery practices, see our Success Playbook on [populating and maintaining your CMDB with discovery](#).

## Service mapping

ServiceNow Service Mapping overlays service maps onto existing configuration data to connect CIs that underlie a given service, identifying both dependencies among CIs and the service impact of failures at the CI level. Figure 1 shows an example of dependency mapping for a specific CI.

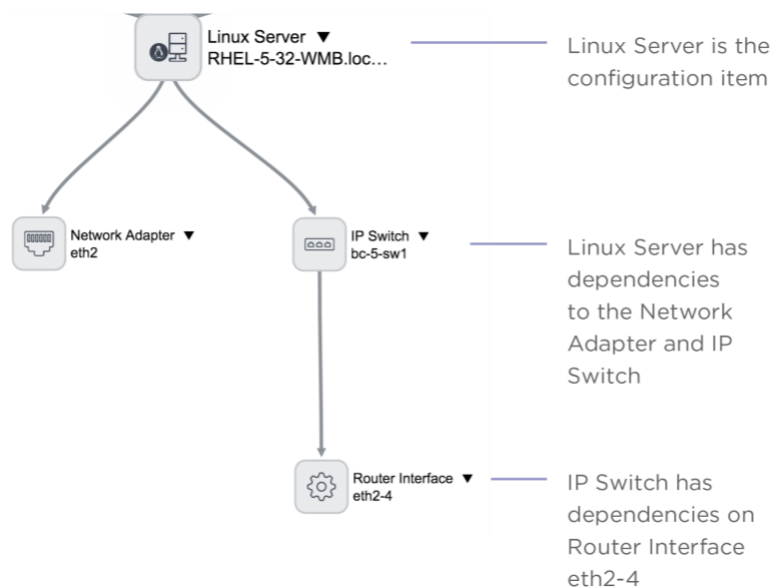


Figure 1: Discovery and dependency map for a specified CI

To get to clear dependency mapping—where your organization can understand the relationship between CIs and specific business outcomes—you must define your organization's business services in Service Mapping. A business service is an abstraction of CIs grouped together in a specific way or for a specific reason to deliver a business outcome. Services can be mapped in one of three ways:

- **Discovered business services** – At a minimum, this requires the organization to define both a name for the business service as well as an entry point—that is, how a business service is consumed by an end user, such as through a URL. The CMDB and Discovery map the interrelated CIs automatically, including an impact tree to show the CI properties, active or related alerts, and outage severity.

- **Manual business services** – Alternatively, you can manually select CIs to include in a business service.
- **Technical services** – Discover dynamic groupings of CIs based on some common criteria, such as the location of all assets of a particular type.

Clear dependency mapping provides the ability to guide effective incident and change management within a service, but not necessarily across services. This requires the organization to identify the relative criticality of services based on the degree they're essential to business operations. If you don't have a well-defined business service taxonomy, you can use your existing disaster recovery and business continuity planning frameworks to understand relative criticality across systems, as defined by recovery tiers.

---

**EXPERT TIP**

Scope the detail you put in your CMDB to what's essential for effective incident and change management automation. Look at historical incident and change data and consult with SMEs to discover the details you need to know about specific configuration items to resolve an incident or understand change risk.

---

## Stage 2 – Build accurate monitoring and event data

### KEY INSIGHTS

- Effective automated incident management depends on the quality of monitoring and event data used to detect and respond to incidents.
- Filter critical from noncritical information with automation so staff can focus on remediation.

Effective incident management automation begins with the ability to separate the “signal”—the monitoring and event data that points to potential disruptions in your business services—from the “noise”—the alerts that reflect noncritical information about the state of your services.

To separate the signal from noise, you should implement a filtering process using ServiceNow [Event Management](#). The steps for this process are outlined below but you can also follow them using the guided setup within the tool.

- 1. Configure a MID Server to receive and process events** – The MID Server (for management, instrumentation, and discovery) is a Java application that runs as a Windows service or UNIX daemon on a server in your local network. It facilitates communication and moving data between your ServiceNow instance and external applications, data sources, and services, including your sources of alert data.
- 2. Configure connector definitions and connector instances to receive external events** – “Connector definitions” specify the MID Server script that pulls events from the external event source.
- 3. Configure event field mappings and alert binding to manage alert generation** – Event field mappings are rules that are used to map values from specific fields to values in other fields. These rules apply after event rule processing and just prior to alert generation, for example, to map event severity fields from a monitoring tool into your ServiceNow severity values. Alert binding automatically binds alerts to CI information from the CMDB. When these two things occur together, they ensure that the alert data is both consistent and clearly maps to CIs.

For more information on these steps, see [Event Management setup](#).

When you complete these steps together, there's less event noise generated by third-party monitoring tools, and you create actionable alerts to help your IT organization resolve service outages.

Events are processed through filters (via the MID Server) that normalize and deduplicate incoming event streams that generate alerts, reducing noise by up to 99%. You can set this up

for discovered business services, manually defined business services, technical services, and alert groups.

When an event from an external source is identified, Event Management locates the CI information to generate an alert, per step 3 above. This CI information is stored in the CMDB through Service Mapping, Discovery, manual entry, and third-party sources.

Service Mapping provides the ability to correlate alerts to relative service impact—and if you have enabled Service Analytics, you'll find additional correlated alert group and root-cause analysis information to help you drive remediation and resolution. Figure 2 depicts the Event Management process flow.

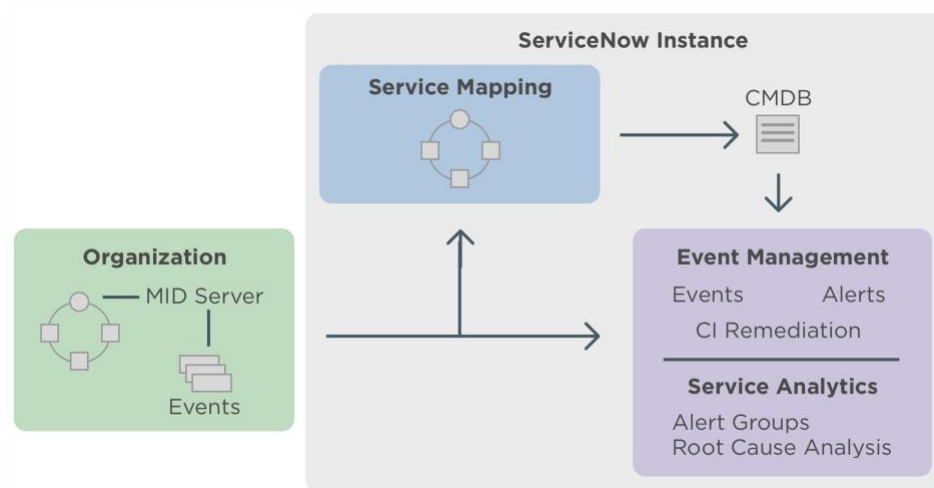


Figure 2: Event Management process flow

Once configured, ServiceNow Event Management enables IT operations teams to view the impacted services and related alerts in a single console, like the one shown in Figure 3. You can select a service in the dashboard filters to show only relevant alerts, or you can select an alert to highlight the impacted services.

You can also view services based on their business criticality, severity, and cost—this helps with prioritizing your remediation and resolution efforts. When you drill into a service, you can identify the probable cause of an impact simply by looking at it.

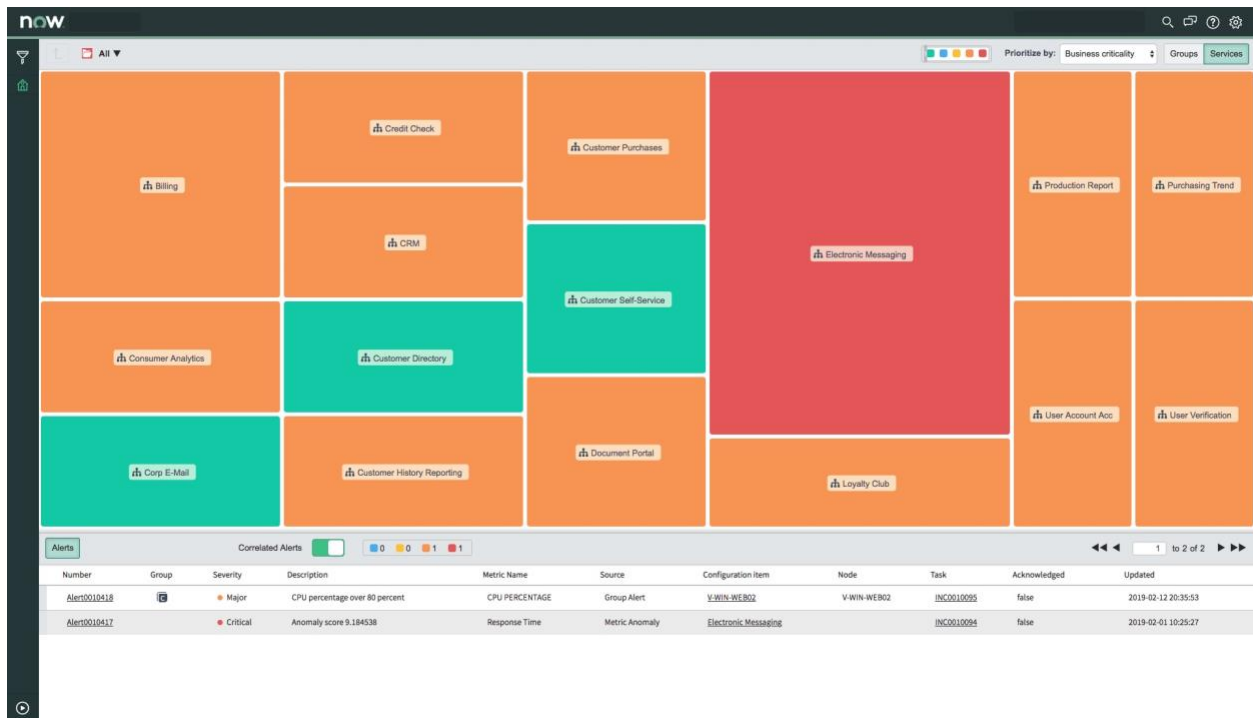


Figure 3: ServiceNow Event Management console

Apply rules to alerts to trigger incident management workflows, including rules to:

- **Autogenerate and assign high-priority incidents based on severity** – This requires clear prioritization and escalation rules, as described in Stage 4.
- **Associate alerts with relevant knowledge base articles to support resolution** – To do this, you need a process for effective knowledge base maintenance, as described in Stage 4.

You can also define alert rules to present automated remediation options through integration with ServiceNow Orchestration.

In all cases, base your alert trigger automation on a clear understanding of how your incidents are prioritized across services, how incidents should be optimally assigned and escalated, and how incidents should be remediated ideally based on historical data.

#### EXPERT TIP

Use subject matter experts—typically technology asset owners—to define a set of remediation options you can present to service desk staff for the most common incidents. This will reduce the time from alert notification to response and resolution.



## Stage 3 – Automate incident response and resolution

### KEY INSIGHT

- For effective incident management automation, get clarity on how the organization defines services, on business rules for how service incidents are prioritized, categorized, and assigned, and on processes for escalation, notification, and knowledge management.

Effective incident management automation requires three things:

- Consolidated real-time information about services and their event status
- Clear prioritization rules based on business impact
- Clear triggers for notifying and escalating incidents to support teams and fast access to the “right knowledge” for support teams and service customers

### Consolidated real-time information about services and their event status

At a minimum, you need an organizational framework for business services, whether these services are discovered through Service Mapping or defined manually from the top down. If you still need some help in this area, go back to Stage 2.

The services reflected in the Event Management console shown in Figure 3 need to reflect services as the organization understands them—such that when the service desk receives an incident notification, it doesn't require significant “translation” to understand where the incident might match to event and alert data.

One possible test is to look at historical data for service outages and determine whether the service, as your support teams define and understand it, clearly matches to the service defined in the Event Management console.

### Clear prioritization rules based on business impact

Support teams can create rules to autogenerate high-priority incidents for alerts based on severity, as well as trigger service desk workflows and automated remediation. To do this, you must have clear, well-defined service level agreements (SLAs) that guide service outage prioritization based on their business impact.

As mentioned in Stage 1, your organization must identify the relative criticality of its services, based on the degree that they're essential to business operations.

## Clear triggers for notification and escalation among support teams and fast access to “right knowledge” for support teams and service customers

Automating incident management isn't effective if it stops at detection and correlation. To improve MTTR, you need to automate your support and remediation activities. Following this checklist will get you there:

- **Be clear about how you should assign different incident categories and types** – To automate your support team activities, begin with routing incidents to the appropriate resolution group. Use historical incident data, along with simple tabletop exercises, to help you define the routing rules you can automate with ServiceNow.
- **Create clear rules to govern escalation and notification** – Support your routing with clear, consistent guidelines for incident escalation—escalations based on business impact, the affected CI(s), or other criteria—and bidirectional notification so escalations can be acknowledged. Again, use historical data related to incidents and service outages, along with simple tabletop exercises, to sharpen the escalation criteria you can automate using ServiceNow On-Call Scheduling.
- **Create a process to maintain your knowledge base** – Ideally, the rules applied to alerts should associate incidents with the appropriate knowledge base article to give support staff guidance on issue resolution. But the effectiveness of these rules depends on how successful your knowledge base maintenance process is. At a minimum, your support teams should follow a process to use your knowledge base to document known errors and fixes. More proactive organizations provide incentives for support teams to rate knowledge base articles, provide feedback, and update and/or contribute articles as part of their daily work.
- **Create channels for effective team collaboration** – Getting access to the “right knowledge” is as much a matter of providing support team collaboration channels as it is of having a populated knowledge base. To maintain the integrity of the incident management process, make sure collaboration takes place in the same system of record as the incident management workflow. To do this, launch chat rooms using On-Call Scheduling, which can log support team conversations, and then channel your collaborations to the chat rooms.

For service customers, effective automation requires three things:

- An omnichannel ability to log incidents—like self-service, chat, email, phone, etc.—both from a desktop and mobile environment
- Access to knowledge base articles that can support self-service resolution (For more on this, see our Success Playbook on [how to improve self-service](#).)
- For business managers, access to real-time reporting on the status of business-impacting service outages

Ensure that service customers are notified of the actions you've taken and the progress you've made toward resolving the incident. Make sure your service desk has clear business rules and standard communication guidance on notification.

---

**EXPERT TIP**

Use tabletop exercises with service desk staff to identify gaps in business rules governing incident management workflow, or identify opportunities for improved collaboration and knowledge sharing.

---

## Stage 4 – Automate change management

### KEY INSIGHTS

- Begin your change management automation strategy with clear definitions of business rules for change risk assessment, scheduling, and approvals and oversight.
- Gain a clear understanding of risk, backed by historical data, so you can expedite more standard, low-risk changes with fewer manual steps.
- Use Change Approval Policies to configure and automate approval governance, expediting normal and emergency changes.
- Use the REST API suite to integrate Change Management with your DevOps teams' continuous integration/continuous delivery pipeline.

With change management automation, your organization can speed up the deployment of changes, reduce the risks associated with change, and support continuous integration/continuous delivery (CI/CD) pipelines managed by DevOps teams. Focus your automation on the specific change management subprocesses that are designed to mitigate risk but that are still largely dependent on manual work—these will likely include change risk assessment, change scheduling, and change approvals and oversight.

### Change risk assessment

When you automate your change risk assessments, you improve the efficiency of your change management process by expediting change risk categorization. This way, you can expedite more standard, lower-risk changes.

Using the Now Platform, you have two methods for automating risk assessment: the Change Risk Calculator and Change Management Risk Assessment.

### Change Risk Calculator plugin

The Change Risk Calculator is activated by default and uses predefined properties and values to calculate a risk value. The system administrator specifies how and when to apply risk and impact rules to change requests. System admins can do this using a builder that sets these conditions (and the order they're evaluated by), through developing more advanced conditions based on business rules, or through a script.

You can also set this common advanced condition based on business rules: *Determine the impact to business services from one or more CIs*. Set a rule to identify if the CI represents a business service, and if the business criticality associated with that service is **1 – most critical** or **2 –**

**somewhat critical.** If the condition matches, the rule will set the risk for the change request to High and the impact to **1 – high**.

Another common scenario that requires scripting is determining the business services that will be impacted as a result of a change to one or more CIs. The calculator plugin provides a sample rule that grabs the CI entered on a change request and locates all associated parent and child business services. These services are then evaluated to identify any critical services.

### Refresh impacted services

The Refresh Impacted Services UI can populate the Impacted Services/CIs related list with any application service impacted by a configuration item listed in the Affected CIs related list (previously, this only worked against the Primary configuration item referenced on the form). The new Impacted Services/CIs related list is generated from all CIs listed in the affected CIs related list (`change.refresh_impacted.include_affected_cis`).

## Risk Assessment

With Change Management Risk Assessment, you have the option to allow end users to answer risk assessment questions associated with a change request so others can use the information to calculate the risk.

This means your organization needs to define risk assessment questions, thresholds, and conditions to support the calculation. Keep your conditions simple and mutually exclusive so they're easy to understand and maintain.

It also means your organization needs to determine weights that reflect the relative importance of assessment question categories, as well as weights for the questions in the assessment categories. You can use historical data from past changes to inform this weighting—the key is to identify the factors that mattered most in the changes that resulted in service interruptions and/or that had to be backed out.

Use the Change Risk Calculator and Risk Assessment together, or just use one. If you use them together, always select the highest risk value from both methods.

Regardless of the method you choose, for effective automation, you need a clear definition of business criticality, whether you use Service Mapping, as outlined in Stage 1, or use the rules and conditions built into Risk Assessment.

To get that clear definition of business criticality, your organization must identify the relative criticality of services, based on how essential they are to business operations. And don't forget—if you don't have a well-defined business service taxonomy, you can use your existing disaster recovery and business continuity planning frameworks to understand relative criticality across systems, as defined by recovery tiers.

Next, the change management process owner and change advisory board (CAB) should review the risk conditions, their threshold values, and the order they're evaluated by.

Automation effectively “cements” a set of business rules for risk assessment so the assessments can be completed quickly and consistently. This means your business rules need to be reviewed and agreed upon upfront—potentially by your CAB—and regularly revisited as services change or as your organization's risk management posture changes.

#### Checklist for automating change risk assessment

- Build—and order—business rules to determine risk.
- If you use Risk Assessment, define risk assessment categories, questions, weights, thresholds, and conditions.
- Review all assessment factors (business rules, conditions, weights, thresholds, etc.) with the change management process owner and CAB for approval.
- Revisit all assessment factors as the service taxonomy changes or as your organization's risk management posture changes.

## Change scheduling

Change scheduling automation should provide the ability to detect and assess whether a planned change conflicts with other changes or blackout periods, and if the proposed change timing meets with predefined maintenance windows. This requires administrators to use the Blackout Schedule form to specify times during which change requests should not be scheduled, and the Maintenance Schedule to specify times during which change requests should be scheduled.

Conflict detection supports effective automation of scheduling by identifying when changes are scheduled at the same time or impact the same CIs. You can run conflict detection manually based on the proposed start/end dates of the change request, the CI subject to change, and (in advanced mode) related CIs.

Automated conflict detection can then run at specific intervals or when a change request is updated. Prior to running conflict detection, however, your organization should take into account:

- **CMDB list size and relationship complexities** – Conflict detection will take longer to complete in organizations with a large CMDB. Don't let this inhibit you from automating conflict detection—just remember to keep the CMDB manageable.
- **Inactive changes are not evaluated** – Automated conflict detection will not evaluate inactive changes when identifying conflicts, such as those for cancelled or closed change requests. This means that change management process owners should put the appropriate procedures and training in place to ensure that cancellations and closures are recorded in a timely fashion on the Now Platform.
- **Advanced mode conflict checking switched off by default** – After an upgrade, advanced mode conflict checking, which identifies conflicts with CIs that are related to or affected by the CI being changed, is switched off. Use risk assessment results to identify whether you should use advanced mode—you'll likely want to know if changes carrying a higher degree of risk have the potential for change collision based on related items, not just the CIs being changed.

Changes at risk of collision are subsequently flagged for approvers. Starting with ServiceNow's Madrid release, change conflict scheduling can also identify additional dates and times where no conflict exists, once a potential change scheduling conflict is identified. Change requesters and fulfillers can take advantage of the **Scheduling Assistant** user interface to search for and select a date and time with no conflicts. To use this, you need to activate two properties:

- Define the number of days to be factored after the respective planned start/end date of a change record when searching for the next available time (change.conflict.next\_available.schedule\_window).
- Define the number of suggestions to be calculated for the next available time field on a change (change.conflict.next\_available.choice\_limit).

In addition to this, ServiceNow has also introduced a new conflict type called **Assigned to already scheduled** to identify when the assignee of a change request is already assigned to another change request scheduled on the same date/time.

Organizations should use the existing calendar interface in the Change Management application to show when changes are planned and where potential conflict exists with either blackout or maintenance schedules.

### Checklist for automating change scheduling

- Build blackout and maintenance schedules.
- Identify whether you need advanced mode conflict detection based on the risk associated with the change request.
- Run manual conflict detection for new change requests, and set automated conflict detection to update detection based on specific time intervals or changes to the CI being changed.
- Activate the Scheduling Assistant to find alternative dates/times for changes when potential conflicts are identified.
- Review all changes regularly to ensure that cancelled or closed changes are marked as inactive.

## Change approvals and oversight

The Now Platform supports the three types of service changes (or “state models”) described by ITIL. You can add new types, but they may modify default workflows. (Consequently, after adding one, review the default workflows for any changes to the state model). Table 1 outlines the tasks required to implement automation for each state model.

TYPE OF CHANGE	DEFINITION	AUTOMATION TASKS
Normal	Any service change that is not a standard or emergency change	Define review and approval rules – Normal changes are routed first for peer review and technical approval. They’re then routed to the change management process owner and CAB for scheduling and final authorization. To do this, you must designate review and approval authorities.
Standard	A pre-authorized change that is low risk, relatively common, and follows a specified procedure or work instruction  Once approved, standard changes can bypass CAB approval.	Define a catalog of pre-authorized change templates – These templates make accessing and requesting a standard change more efficient. They include rebooting a Windows server, adding a network switch to a data center cabinet, standard implementation, blackouts, and test plans for the change, so new proposals only need to include the CI(s), assignment group, and schedule. Base your templates on a history of similar successful changes. Note that you can track the success of pre-authorized templates—this way, you can modify or retire them if they don’t lead to successful changes.
Emergency	A change that must be implemented as soon as possible, for example, to resolve a major incident or implement a security patch	Define approval rules, which may vary from normal approvals – Emergency changes progress directly from submission to the Authorize state for approval from the relevant CAB approval authority. Some organizations define an emergency CAB (e-CAB), with specific authorities for approving emergency changes.

Table 1: ServiceNow state models for change management

Starting with ServiceNow’s Madrid release, you can use **Change Approval Policies** to automate the generation of change approval actions based on a set of predefined conditions. Change Approval Policies are defined using decision tables—a Now Platform feature that provides the



ability to define the logic for complex decisions dependent on multiple inputs. This allows you to define multiple, variable inputs to your change decision and evaluate them dynamically to generate the appropriate change approvals action.

To use Change Approval Policies, complete these steps:

1. Replace Group Approval activities or User Approval workflow actions with the Change Approval Policy activities in the Normal and/or Emergency change management workflow.
2. Configure approval policies based on three components:
  - a. **Policy inputs** – The inputs that will be evaluated as part of your decision
  - b. **Decisions** – Records that contain the conditions that determine what approval action is to be taken
  - c. **Approval definitions** – The approval action to take when inputs match decision conditions

See Figure 14 for a screen shot showing how Change Approval Policies are configured. *Note that while Change Approval Policies use conditional logic, the logic contained in a policy takes into account all policy inputs before determining an approval action.* This means that a decision does not simply look at inputs sequentially in a standard *if-then-else* construct—rather, it reviews all inputs before matching decision conditions.

The screenshot displays the 'Decision Manager approval' configuration page for 'Low Risk Manager approval'. The interface includes a header with navigation icons and buttons for 'Update' and 'Delete'. Below the header, there is a section for 'Label to use for this Decision condition' with a dropdown menu set to 'Low Risk Manager approval' and an 'Order' field set to '0'. The 'Answer' field is set to 'Change Approval Definition: Assign'. Below this, a section titled 'All of the following conditions must be true, for this answer to be used' contains three conditions: 'Change request.State' is 'Authorize', 'Change request.Risk' is 'Low', and 'Manager approved' is 'false'. Each condition is followed by a dropdown menu with 'AND', 'OR', and 'X' options. At the bottom, there are 'Update' and 'Delete' buttons.

Figure 14: Configuring change approval policies

## Using change approval policies to support DevOps

The conditional logic in Change Approval Policies can take inputs from Test Management to dictate an automated approval or rejection. In this case, a test for a specific change would represent a policy input. Documented completion and success of that test in Test Management would represent the decision conditions for approval, which could be defined as deployment of the change. See [Applying Change Approval Policies to DevOps](#) in the ServiceNow Community for more detail.

A new change request progresses through states once you submit it. Table 2 shows these change states. In the Assess state, ServiceNow routes normal changes to peer reviewers and technical approvers, so you need to identify reviewers and approvers (ideally in advance, using Change Approval Policies) for efficient automation.

CHANGE STATE	DESCRIPTION
New	Change request is not yet submitted for review and authorization. A change requester can save a change request as many times as necessary while building out the details of the change prior to the submission.
Assess	Peer review and technical review of the change details are performed during this state. This is not required for <b>Standard</b> or <b>Emergency</b> changes.
Authorize	Change management and the CAB schedule the change and provide final authorization to proceed. This is not required for <b>Standard</b> changes.
Scheduled	The change is fully scheduled and authorized and is waiting for the planned start date. An email notification is sent to the user who scheduled the change.
Implement	The planned start date has approached and the actual work to implement the change is being conducted. An email notification is sent to the user who requested the change.
Review	The work has been completed. The change requester determines whether the change was successful. A post-implementation review can be conducted during this state. An email notification is sent to the user who requested the change.
Closed	All review work is complete. The change is closed with no further action required.
Canceled	A change can be canceled at any point when it is no longer required. However, a change cannot be canceled from a <b>Closed</b> state. An email notification is sent to the user who requested the change.

Table 2: ServiceNow change states

The ServiceNow CAB Workbench plugin supports streamlined scheduling automation and CAB approvals. To put the plugin to use, the change management process owner must ensure that the organization has:

- A clearly defined CAB manager
- A list of CAB members (including those who can substitute for the CAB manager) for normal and emergency CAB meetings
- A schedule for CAB meeting dates
- Clear rules on the types of change requests that are to be included in the CAB agenda

Once established, your CAB members can view an overall meeting calendar and individual agenda items (i.e., change requests) alongside a change calendar that includes the blackout schedule and maintenance schedule.

CAB meeting chairs can also align individual agenda items to specific time slots, so stakeholders only need to attend a CAB meeting based on automated notification of when discussion starts for relevant agenda items.

The value of automating and overseeing your change approvals is largely driven by the extent that change management process owners can shift overall change volumes from normal to standard changes. Change management process owners should provide incentives and KPIs for teams to build pre-authorized change templates that users can access through a change catalog. Teams should target CIs with a consistent, successful, and frequent record of change that's based on historical data and/or interviews with subject matter experts. Ideally, standard changes should constitute 70–90% of total change volumes.

---

**EXPERT TIP**

If your organization wants to implement DevOps and/or continuous delivery, focus your effort upfront on building a change catalog of pre-authorized templates for the CIs anticipated to change most frequently.

---

## Implementing the REST API suite to support CI/CD

The Change Management API available with ServiceNow's Madrid release provides REST APIs that enable third-party application integration with ServiceNow Change Management. The suite of APIs allows teams to create, update, approve, and work on changes from the CI/CD pipeline from change creation to closure. Specifically, the Change Management API enables integrators to:

- Initiate standard changes from a published standard change request template
- Create normal or emergency change requests
- Update fields in the change request table, task table, and work tasks from creation to closure
- Retrieve change requests, standard templates, or change request tasks
- Generate and process approval activities
- Identify potential scheduling conflicts
- Delete change requests, change request tasks, and conflict checking processes

Work with integration and DevOps teams to build this functionality into your CI/CD toolchain, to reduce the need for “swivel chair” behavior between systems and ensure a single system of record for changes.

## Stage 5 – Measure impact and tune

### KEY INSIGHTS

- Use Performance Analytics to find ways to improve your automation.
- Use a short list of questions to select the KPIs needed to tune your automation strategy.

Effective automation is not a one-time project but an ongoing process that looks for new opportunities and identifies where automation is (or isn't) delivering value. ServiceNow Performance Analytics provides access to several out-of-the-box KPIs that can help you measure opportunities to improve and tune your incident and change management automation strategy. See Tables 3 and 4 for incident and change management KPIs and how you can use them to tune your automation.

KPIs	HOW TO TUNE YOUR AUTOMATION STRATEGY
<ul style="list-style-type: none"> <li>• Average age of the last update of open incidents</li> <li>• Percentage of open incidents not updated in the last 30 days</li> <li>• Number of open incidents not updated in the last 5 days</li> <li>• Number of open incidents not updated in the last 30 days</li> <li>• Summed age of last update of open incidents</li> <li>• Percentage of open overdue incidents</li> <li>• Number of open overdue incidents</li> </ul>	Review your business rules (and team KPIs) for incident notification, routing, and escalation.
<ul style="list-style-type: none"> <li>• Percentage of incidents closed by self-service</li> <li>• Number of incidents closed by self-service</li> <li>• New/open workloads</li> <li>• Incident backlog growth</li> </ul>	Ensure your knowledge base reflects the most common self-service use cases for end users and provides clear guidance.
<ul style="list-style-type: none"> <li>• Percentage of open incidents reassigned at least once</li> <li>• Average reassignment of open incidents</li> <li>• Summed reassignment of open incidents</li> <li>• Percentage of incidents resolved without reassignment</li> <li>• Number of reassigned open incidents</li> <li>• Number of open incidents unassigned</li> </ul>	Review the business rules (and team KPIs) for incident routing and escalation.
<ul style="list-style-type: none"> <li>• Number of incidents not solved</li> <li>• Percentage of incidents not solved</li> </ul>	Ensure the knowledge base reflects the most current guidance for known errors.
<ul style="list-style-type: none"> <li>• Average resolution time of resolved incidents</li> <li>• Average close time of incidents</li> </ul>	Identify primary drivers for incident cycle time: A lack of current guidance in the knowledge base may require unnecessary escalation, or your business rules for notification, routing, and escalation may require updating to reflect the fastest path to resolution.

Table 3: Selected incident management KPIs (available in ServiceNow Performance Analytics) and suggested actions to tune incident management automation

KEY PERFORMANCE INDICATORS	WHAT TO TUNE IN YOUR AUTOMATION STRATEGY
<ul style="list-style-type: none"> <li>Change backlog growth</li> </ul>	Build pre-authorized templates to increase the availability of standard changes in the change catalog.
<ul style="list-style-type: none"> <li>Number of unsuccessful changes</li> </ul>	Reevaluate risk assessment: Determine whether you need to make adjustments to business rules in the Change Risk Calculator or to the risk assessment questions for change requesters.
<ul style="list-style-type: none"> <li>Number of reassigned open changes</li> <li>Percentage of open changes reassigned at least once</li> <li>Average reassignment of open changes</li> <li>Summed reassignment of open changes</li> </ul>	Review change request assignment rules for validity.
<ul style="list-style-type: none"> <li>Average age of open changes</li> <li>Summed age of open changes</li> <li>Average age of "updated since of" open changes</li> <li>Summed age of "updated since of" open changes</li> </ul>	<p>Build pre-authorized templates to increase the availability of standard changes in the change catalog.</p> <p>Identify the change state (see Table 1) where open changes spend the most time. Based on the change state, determine whether you need to make changes to the business rules for change risk assessment, change request assignment, or approvals and scheduling.</p>

*Table 4: Selected change management KPIs (available in ServiceNow Performance Analytics) and suggested actions to tune change management automation*

You don't have to use all of the KPIs reflected in Tables 3 and 4—and there are more KPIs that you can use to track the performance of incident and change management. What's most important is that you select and use KPIs that can provide directional guidance about where you can improve your automation. If a KPI begins to trend negatively, the process owner and team should know how to respond.

Here's an example: Your organization may decide to monitor change backlog growth as an indicator of the efficiency of your change management process. If the change backlog grows—potentially slowing down necessary maintenance and enhancement activities—the team should determine if the backlog growth is the result of a manual slowdown in the overall process. One potential root cause is an increase in the ratio of normal changes—those requiring CAB approval—from standard changes with pre-authorized approval. In this case, the team should look to increase the number of standard changes made available in a change catalog, through the development of pre-authorized templates, for changes with a successful track record.

To identify the right KPIs to guide automation tuning, ask the following questions:

- **What business rules have we put in place to guide our automation?** For incident management, this should include how incidents are categorized, as well as the rules for routing, escalation, and notification. For change management, this should include rules governing risk assessment, routing and approvals, and scheduling.
- **What knowledge have we made available as part of our automation plan?** This should include your knowledge base (for service desk agents as well as end users) and a change catalog for standard changes.
- **What changes in performance do we expect to see as the result of our automation plan?** Teams should ask how a business rule, set of business rules, or publication of knowledge is intended to change performance. This should relate to your overall goals for the process and

organization. For example, a team may design business rules for incident management prioritization such that high-severity incidents are autogenerated based on alerts for business-critical services. The intent behind this is to speed incident remediation for business-critical services, meaning that the team should look for reductions in the MTTR incidents for business-critical services.

- **If we don't see the change we want, what do we need to tune?** If the team aligns the right KPIs to business rules and publication of knowledge, you then have clarity on where to look first to improve your automation. In the example outlined above, let's assume that the MTTR for incidents affecting business-critical services refused to budge—or, worse, got longer. The team should first investigate whether automation is working—was the rule that autogenerates alerts for business-critical services written correctly? Some business-critical services may not have been identified correctly in the CMDB, or alert generation may not be triggering the right service desk workflows. No KPI provides an immediate solution, but it can guide the team on where to look first to find one.

---

**EXPERT TIP**

More metrics does not equal better metrics. Select a small set of KPIs that your process owners and teams readily understand and that give consistent directional guidance on process improvement. If teams don't know what to do to move a KPI, then it needs reevaluation.

---