

ITIL® EVENT MANAGEMENT



[ITIL 4 Guide >](#)

(This article is part of our [ITIL v3 Guide](#). Use the right-hand menu to navigate.)

ITIL event management

Thousands (or millions) of events happen across your IT infrastructure every day. In large enterprises, the number could be billions. Why? Because an event is simply a change to the state of an IT service or configuration item (CI) that is significant to its management.

A server moving from online to idle could be an event, or the completion of a regular server maintenance script: they're worth knowing about, and there may even be an action you wish to take as a result.

The objective of event management is to detect events, analyze them, and determine the right control action (if any). By doing so, the event management process also provides a strong foundation for service assurance, reporting, and service improvement.

It's important to know, though, that monitoring and event management are not the same thing. Monitoring is certainly a component of event management, in that it is a useful way to detect events as they occur. Event management, on the other hand, is focused on extracting meaning out of events, to help IT take appropriate actions (when required).

Download Now: ITIL 4 Best Practice e-Books

These all-new for 2020 ITIL e-books highlight important elements of ITIL 4 best practices. Quickly understand key changes and actionable concepts, written by ITIL 4 contributors.

[Free Download >](#)



[Free Download >](#)

The scope and benefits of event management

Event management can be applied to any aspect of service management that needs to be controlled and which can be automated — from networks, servers, and applications all the way to environmental conditions like fire and smoke detection and [security](#) and intrusion detection.

Since event management can be applied to just about every aspect of service management in your IT organization, the benefits are widespread. In general, effective event management practices can:

- Provide a strong foundation to automate key components of your IT operation
- Improve detection and response times to incidents, changes, exceptions, etc.
- Reduce downtime as a result of the above

So what does success look like? In event management, success is being able to detect, communicate, and take the appropriate action for every event (or change in state) that is significant to managing your IT services and the CIs that support them.

Event management process flow

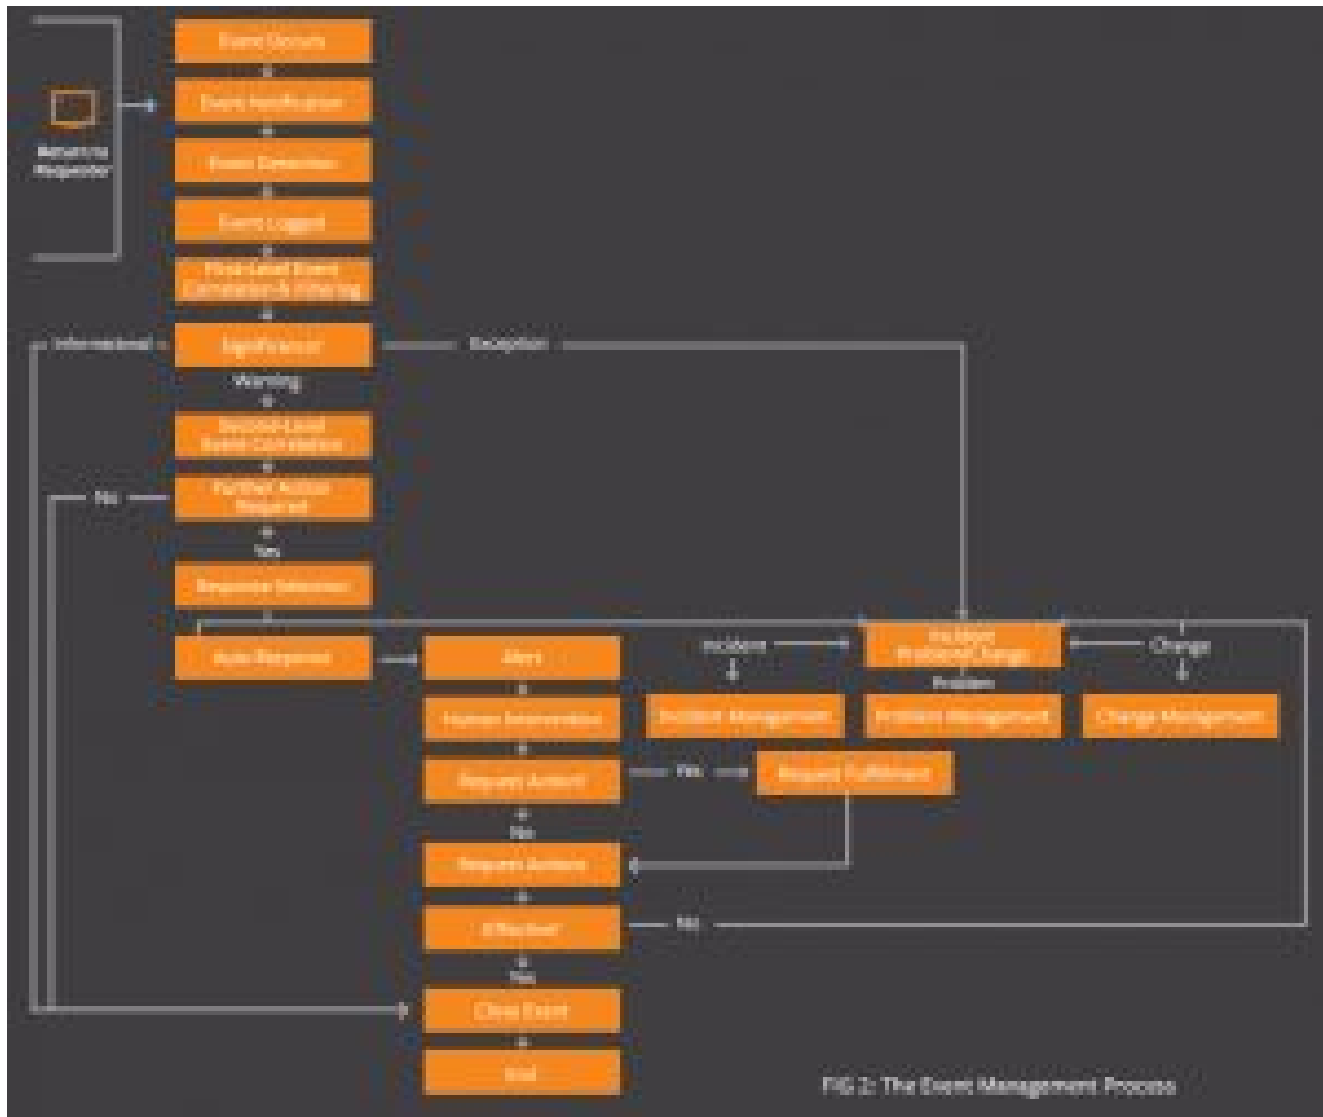


FIG 2: The Event Management Process

What's the difference between events and incidents?

It's a great question, and the answer is simple. Incidents are unplanned interruptions or significant reductions in the quality of an IT service. When an incident occurs, something is wrong. Events, on the other hand, are simply changes in the state of your services, CI's, or pretty much anything of significance across your IT infrastructure.

So can an incident be an event? Absolutely. All incidents are events, since an outage or service quality reduction is a change in the state of that service. But not all events are incidents, since an increase in utilization, a user logging in, or an automated backup service completing represents a status change, but not a disruption or degradation in service quality.

In fact, there are three types of events defined by ITIL:

1. **Information.** These events typically don't require a response of any type, since they are basic status updates, or data generated to aid with reporting, etc. Logs and reports are great examples.
2. **Warning.** Warnings are indicators of activity that outside the norm — like a threshold being approached. Like a hurricane or tornado warning, a warning means that you should monitor conditions to make sure they do not worsen — or take action to prevent them from worsening when appropriate. An example of this type of event would be server capacity reaching 75%, or

a standard transaction taking 15% longer to complete than normal.

3. **Exception.** Exception events are indicators that something is wrong. The services (and business they support) may be negatively impacted. A network or server being down (as opposed to just approaching capacity) is an example of an exception.

What other activities could be considered events and trigger the event management process? Quite a few — from exceptions to automated processes to simple status changes in a server or database. The sky is the limit.

It is ultimately the job of IT to designate what types of activities they will consider information events, warning events, and exception events. As a general rule, though, you will want to categorize an event as “information” when it will purely be used to gain insight and inform better decision-making. “Warning” events are typically those that may require closer monitoring or even intervention to help you prevent exceptions from occurring. “Exception” means something is really wrong that typically requires immediate action.

The key activities of event management

During the design phase of your IT services, you should define which types of events need to be generated, and how they will be generated, for each type of configuration item (CI) involved in delivering the service. The typical event lifecycle is:

1. **Event occurrence**

Events occur 24 x 7 x 365. In ITIL Event Management, the key is defining the types of events that are significant to your operation and ensuring you have a system in place to detect them.

2. **Event notification**

Notifications are typically sent by monitoring tools or CIs (configuration items). At this stage, these are simply notifications that an event has happened — and have typically not yet been interpreted or correlated to understand the meaning or impact.

3. **Event detection**

In this step, a monitoring system, automated agent, or systems management solution receives the notification and determines the meaning of the event.

4. **Event logged**

A record of the event is made, along with any subsequent actions taken. This may be done by your systems management solution, or by the individual applications / services / hardware that triggered the event.

5. **Event filtering and correlation**

Can the event be ignored, or does it need to be passed on to the events management system? Often, information events are ignored. Warnings and exceptions often require additional action, though. So the first step of this process — called first-level correlation and filtering — is simply filtering which events should be ignored versus passed on to the event management system. In the second level of correlation, a correlation engine uses predefined business rules to determine the significance of warning and exception events, and decide the appropriate next steps.

6. **Event response / further action**

Remember, all events (and responses) should be logged. In addition, based on the event type and severity, the correlation engine may determine it is appropriate to escalate the event to a team or individual, or in the case of more severe warnings and exceptions, even automatically

create an incident, problem, or change.

7. Closing the event

If an event results in an incident, problem, or change being created, event closure should be handled through those respective processes. They can be “closed” in the event management system by ensuring the event is properly logged as well as the subsequent action taken, and including a link to the corresponding incident, problem, or change request. Like most other ITIL process, event management doesn't live in a bubble. While event management primarily interfaces with incident, problem, and change management (for dealing with exceptions), it also interfaces with:

- Capacity and availability management for understanding the significance of events, thresholds, etc.
- Asset Management for managing the status of assets
- Configuration Management, for managing the status of CIs.

Measuring Your effectiveness

To help you gauge the efficiency and effectiveness of your Event Management process, these are just a few of the KPIs you can track.

- The number or percentage of events that become incidents.
- The CIs that generate the most events
- How many events are reported by your monitoring tools, and the breakdown by event category
- The total percentage of events that become incidents (or alternately result in changes), and more specifically, how many of these incidents are reported by your automated systems.

Key recommendations

First, be sure to perform a thorough study of the types of events that occur in your IT environment. Know which systems log events, and where, and what the events mean.

That makes it much easier to understand and define which types of events require additional care — whether it's human intervention or automated workflows for handling changes or raising incidents.

Since it's not humanly possible for a live person (or even team of people) to monitor and manage every event triggered by all of your systems, your goal is to create a simple, streamlined set of workflows to automate the easy stuff — and alert your team when more significant events that threaten services (or that require human assistance of any type) occur.

Finally, make sure your event logs are capturing the appropriate level of details — what happened, when it happened, how it was handled, who it was escalated to, and any details of communication with other people or systems to support any actions taken. You'll also want to capture whether events are breaching any of your SLAs or OLAs, to help you remain compliant and provide accurate reporting.