Notification escalations

Introduction

Shinken supports optional escalation of contact notifications for hosts and checks. Escalation of host and check notifications is accomplished by defining escalation objects.

When Are Notifications Escalated?

Notifications are escalated if and only if one or more escalation definitions matches the current notification that is being sent out. If a host or check notification does not have any valid escalation definitions that applies to it, the contact group(s) specified in either the host group or check definition will be used for the notification.

Look at the example below:

Property	Value
Name	To-level-2
first_notification_time	60
last_notification_time	120
contact_groups	nt-admins,managers

It will use the interval length for the value you set for first/last notification time. Here, it will escalate after 1 hour problem, and stop at 2 hours.

Contact Groups

When defining notification escalations, it is important to keep in mind that any contact groups that were members of "lower" escalations (i.e. those with lower notification time ranges) should also be included in "higher" escalation definitions. This should be done to ensure that anyone who gets notified of a problem continues to get notified as the problem is escalated.

Example:

Property	Value
Name	To-level-2
first_notification_time	60
last_notification_time	120
contact_groups	nt-admins,managers

Property	Value
Name	To-everyone
first_notification_time	120
last_notification_time	240
contact_groups	nt-admins,managers,everyone

The first (or "lowest") escalation level includes both the nt-admins and managers contact groups. The last (or "highest") escalation level includes the nt-admins, managers, and everyone contact groups.

Notice that the nt-admins contact group is included in both escalation definitions. This is done so that they continue to get paged if there are still problems after the first two check notifications are sent out. The managers contact group first appears in the "lower" escalation definition - they are first notified when the third problem notification gets sent out. We want the managers group to continue to be notified if the problem continues past five notifications, so they are also included in the "higher" escalation definition.

Overlapping Escalation Ranges

Notification escalation definitions can have notification ranges that overlap. Take the following example:

Property	Value
Name	To-level-2
first_notification_time	60
last_notification_time	240
contact_groups	nt-admins,managers

Property	Value
Name	To-everyone
first_notification_time	120
last_notification_time	0
contact_groups	on-call-support

In the example above:

- The nt-admins and managers contact groups get notified on between 1 and 2 hours
- All three contact groups get notified between 2 and 4 hours
- Only the on-call-support contact group gets notified after 4 hours

Escalations based on time short time

It's also interesting to see that with escalation based on time, if the notification interval is longer than the next escalation time, it's this last value that will be taken into account.

Let take an example :

Host:

Property	Value
Name	srv-important
notification interval	1440
escalations	To-level-2

Then with the escalations object:

Property	Value
Name	To-level-2

first_notification_time	60
last_notification_time	120
contact groups	level2

Here let say you have a problem HARD on the check at t=0. It will notify the host contacts. The next notification should be at t=1440 minutes, so tomorrow. It's ok for classic notifications, but not for escalated ones.

Here, at t=60 minutes, the escalation will raise, you will notify the level2 contact group.

So you can put large notification_interval and still have quick escalations times.

Time Period Restrictions

Under normal circumstances, escalations can be used at any time that a notification could normally be sent out for the host or check.

This "notification time window" is determined by the notification period directive in the host or check configuration.

You can optionally restrict escalations so that they are only used during specific time periods by using the "escalation_period" directive in the escalation configuration.

If you use the "escalation_period" directive to specify a time period which the escalation can be used, the escalation will only be used during that time. If you do not specify any escalation period directive, the escalation can be used at any time within the "notification time window" for the host or check.

Escalated notifications are still subject to the normal time restrictions imposed by the "notification_period" directive in a host or check definition, so the timeperiod you specify in an escalation definition should be a subset of that larger "notification time window".

State Restrictions

If you would like to restrict the escalation definition so that it is only used when the host or check is in a particular state, you can use the scalation options directive in the escalation definition. If you do not use the "escalation_options" directive, the escalation can be used when the host or check is in any state.