

Scheduled Downtimes

Introduction

Shinken allows you to schedule periods of planned downtime for hosts and check that you're monitoring. This is useful in the event that you actually know you're going to be taking a server down for an upgrade, etc.

Scheduling Downtime

You can schedule downtime with your favorite UI or as an external command in cli.

Once you schedule downtime for a host or check, Shinken will add a comment to that host/check indicating that it is scheduled for downtime during the period of time you indicated. When that period of downtime passes, Shinken will automatically delete the comment that it added. Nice, huh?

How Scheduled Downtime Affects Notifications

When a host or check is in a period of scheduled downtime, Shinken will not allow normal notifications to be sent out for the host or check. However, a "DOWNTIMESTART" notification will get sent out for the host or check, which will serve to put any admins on notice that they won't receive upcoming problem alerts.

When the scheduled downtime is over, Shinken will allow normal notifications to be sent out for the host or check again. A "DOWNTIMEEND" notification will get sent out notifying admins that the scheduled downtime is over, and they will start receiving normal alerts again.

If the scheduled downtime is cancelled prematurely (before it expires), a "DOWNTIMECANCELLED" notification will get sent out to the appropriate admins.

Overlapping Scheduled Downtime

I like to refer to this as the "Oh crap, its not working" syndrome. You know what I'm talking about. You take a server down to perform a "routine" hardware upgrade, only to later realize that the OS drivers aren't working, the RAID array blew up, or the drive imaging failed and left your original disks useless to the world. Moral of the story is that any routine work on a server is quite likely to take three or four times as long as you had originally planned...

Let's take the following scenario:

- You schedule downtime for host A from 7:30pm-9:30pm on a Monday
- You bring the server down about 7:45pm Monday evening to start a hard drive upgrade
- After wasting an hour and a half battling with SCSI errors and driver incompatibilities, you finally get the machine to boot up
- At 9:15 you realize that one of your partitions is either hosed or doesn't seem to exist anywhere on the drive
- Knowing you're in for a long night, you go back and schedule additional downtime for host A from 9:20pm Monday evening to 1:30am Tuesday Morning.

If you schedule overlapping periods of downtime for a host or check (in this case the periods were 7:40pm-9:30pm and 9:20pm-1:30am), Shinken will wait until the last period of scheduled downtime is over before it allows notifications to be sent out for that host or check. In this example notifications would be suppressed for host A until 1:30am Tuesday morning.