

Active and passive checks

Active Checks

Introduction

Shinken Enterprise is able to monitor hosts and checks in two different ways: actively and passively. Using Active checks is the most common method for monitoring hosts and checks. The main features of active checks are the following:

- Active checks are initiated by the Shinken Enterprise poller process
- Active checks run on a regular scheduled basis

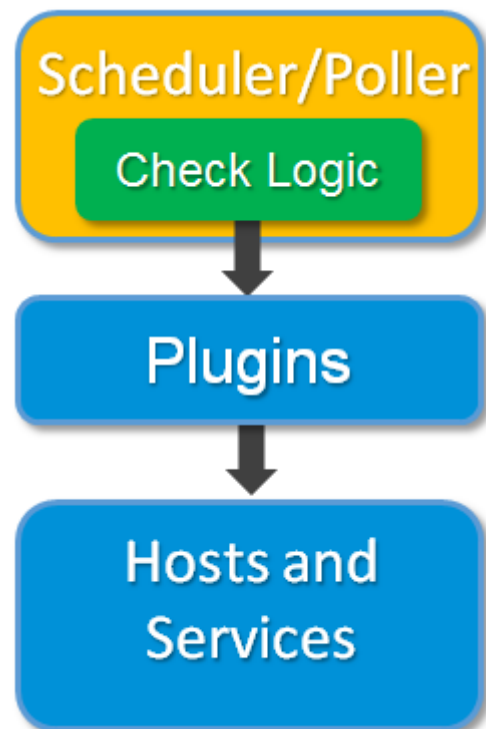
How are active Checks performed?

Active checks are initiated by the check logic in the Shinken Enterprise daemon.

When Shinken Enterprise needs to check the status of a host or a check, it will execute a plugin and show the information about what needs to be checked.

The plugin will then check the operational state of the host or check and report the results back to the Shinken Enterprise daemon.

The scheduler daemon will process the results of the host or check and take the appropriate action if necessary (e.g. send notifications, ask for event handlers, etc).



When Are Active Checks Executed?

Active check are executed:

- At regular intervals, as defined by the "check_interval" and "retry_interval" options in your host and check definitions
- On-demand as needed

Regularly scheduled checks occur at intervals equaling either the "check_interval" or the "retry_interval" in your host or check definitions, depending on what type of state the host or check is in. If a host or check is in a HARD state, it will be actively checked at intervals equal to the "check_interval" option. If it is in a SOFT state, it will be checked at intervals equal to the retry_interval option.

On-demand checks are performed whenever Shinken Enterprise sees a need to obtain the latest status information about a particular host or check. For example, when Shinken Enterprise is determining the reachability of a host, it will often perform on-demand checks of parent and child hosts to accurately determine the status of a particular network segment.

Passive checks

Introduction

In most cases you'll use Shinken Enterprise to monitor your hosts and checks using regularly scheduled active checks. Active checks can be used to "poll" a device or check for status information as often as needed. Shinken Enterprise also supports a way to monitor hosts and checks passively instead of actively. The key features of passive checks are the following:

- Passive checks are initiated and performed by external applications/processes
- Passive check results are submitted to Shinken for processing

The major difference between active and passive checks is that active checks are initiated and performed by Shinken Enterprise, while passive checks are performed by external applications.

Uses For Passive Checks

Passive checks are useful for monitoring checks that are:

- Asynchronous by nature, they cannot or would not be monitored effectively by polling their status on a regularly scheduled basis, like backup or batchs
- Located behind a firewall and cannot be checked actively from the monitoring host

Examples of asynchronous checks that lend themselves to being monitored passively include:

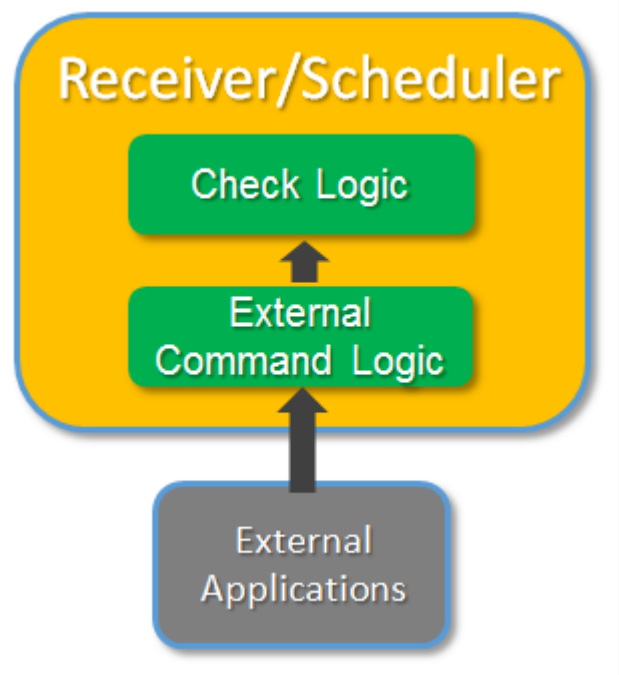
- "SNMP" traps and security alerts. You never know how many (if any) traps or alerts you'll receive in a given time frame, so it's not possible to just monitor their status every few minutes.
- Aggregated checks from a host running an agent. Checks may be run at much lower intervals on hosts running an agent.
- Submitting check results that happen directly within an application without using an intermediate log file(syslog, event log, etc.).

How Passive Checks Work

Here's how passive checks work in more detail.

- An external application checks the status of a host or check.
- The external application writes the results of the check to the webservice of the receiver. Its configuration and API is defined at [Enable webservice for passive checks](#).
- Shinken Enterprise reads the passive checks and push them to the appropriate daemons.
- Shinken Enterprise will get results each second and scan the check result queue. Each check result that is found in the queue is processed in the same way - regardless whether the check was active or passive. Shinken Enterprise may send out notifications, log alerts, etc. depending on the check result information.

The processing of active and passive check results is essentially the same. This allows for seamless integration of status information from external applications with Shinken Enterprise.



Enabling Passive Checks

In order to enable passive checks in Shinken Enterprise, you'll need to do the following:

- Set `accept_passive_service_checks` directive is set to 1 (in `shinken.cfg`).
- Set the `passive_checks_enable` directive in your host and check definitions is set to `True`.

If you want to disable processing of passive checks on a global basis, set the `accept_passive_check_checks` directive to 0.

If you want to disable passive checks for just a few hosts or checks, use the `"passive_checks_enabled"` directive in the host and/or check definitions to do so.

Submitting Passive check Results

You can look at the [Enable webservice for passive checks](#) about how to send external checks to the receivers.

Passive Checks and Host States

Unlike with active host checks, Shinken Enterprise does not (by default) attempt to determine whether or host is `DOWN` or `UNREACHABLE` with passive checks. Rather, Shinken Enterprise takes the passive check result to be the actual state the host is in and doesn't try to determine the hosts' actual state using the reachability logic. This can cause problems if you are submitting passive checks from a remote host.

Passive host checks are treated as `HARD` states.

