

Basic - Hosts and checks

Hosts

In Shinken Enterprise, the basic element of supervision is the host.

- A host represents any object with a network address.
- A host is not limited to physical machines, but can also be a virtual machine, a connected object or any object that is accessible and identifiable on a network.

Shinken Enterprise performs checks on each host at regular time intervals, as well as on the checks attached to each host..

- The control performed on each host allows to verify that the host is reachable by their network address, but may be more complex depending on what has been configured by your Shinken administrator.
- This verification is of course configurable by the person(s) in charge of setting up the supervision.

Host Groups

Guests can be organized in groups to facilitate handling and packing.

One can imagine for example a "Web" host group that would contain all the servers dedicated to Web hosting, to facilitate their handling and improve readability..

Checks

In addition to the default verification, a set of additional and more precise checks can be performed on each host. This role is fulfilled by the checks.

A check represents a particular verification that will be performed on the host to which it is attached. Multiple checks can be attached to the same host, and a check can be attached to several different hosts.



Exemple

There is a dedicated web server to make a website available to the public. Host verification allows us to know if the server can be reached by its address on the network. But we would like to hang checks on it to get more information about its operating status.:

- A check to see if the home page of the site is available,
- A check to see if the server's performance in terms of speed is acceptable,
- A check to see how many users are visiting the site.,

These are only examples, but the set of possible checks on a host is not limited and can be extended as needed.

When verified, a check provides a result and a long result, a status, a context and performance data (optional).

Status and context

Once performed, a check first returns a status and a context.

The status (Critical, Warning, OK, Unknown) as well as the context of the check (DOWNTIME, ACKNOWLEDGED, FLAPPING) indicate the status of the check verification.

The statuses and contexts are described in more detail in the page dedicated to them: [Concept: Status & Context](#)

Result and long result


The verification of a check also provides a result and a long result. It is a textual information that provides the detailed information that the check can provide.

This information is separated in 2 parts:

- The result: This is a brief summary to capture the main information returned by the checklist.
- The long result: More detailed information allowing to have more precise information on the execution of the check. The long result is optional and often absent.

Exemple

The check "Poller - Performance", which allows to supervise the good functioning of Shinken, has a Result and a Long Result.

Statut	Nom de check	Résultat	Résultat Long																														
Pas de sélection	Saisir un nom de check	Saisir du texte	Saisir du texte																														
	Poller - Performance	<p>[OK] Poller statistics :</p> <ul style="list-style-type: none">[2,58] Checks done per second.[65%] normal] Average server RAM usage.[Resources available] Poller load.[4%] Average CPU running of the checks on 4 core.[8] Processes in the CPU Queue. normal	<p>Top 5 checks :</p> <table border="1"><thead><tr><th>Host</th><th>Check</th><th>Cpu ti</th></tr></thead><tbody><tr><td>srv-mysql-2</td><td>srv-mysql-2</td><td>189m</td></tr><tr><td>srv-web-1</td><td>Kernel Stats</td><td>180m</td></tr><tr><td>SERVER1_CUSTOMER2</td><td>Kernel Stats</td><td>179m</td></tr><tr><td>ht1</td><td>SSH Connexion</td><td>170m</td></tr><tr><td>SERVER1_CUSTOMER2</td><td>DiskStats</td><td>160m</td></tr></tbody></table> <p>Nb Checks per CPU running time :</p> <table border="1"><thead><tr><th>0 - 50ms</th><th>50 - 100ms</th><th>100 - 200ms</th><th>200 - 300ms</th><th>300 - 400ms</th><th>+</th></tr></thead><tbody><tr><td>182 (35%)</td><td>36 (6%)</td><td>301 (57%)</td><td>0 (0%)</td><td>0 (0%)</td><td></td></tr></tbody></table>	Host	Check	Cpu ti	srv-mysql-2	srv-mysql-2	189m	srv-web-1	Kernel Stats	180m	SERVER1_CUSTOMER2	Kernel Stats	179m	ht1	SSH Connexion	170m	SERVER1_CUSTOMER2	DiskStats	160m	0 - 50ms	50 - 100ms	100 - 200ms	200 - 300ms	300 - 400ms	+	182 (35%)	36 (6%)	301 (57%)	0 (0%)	0 (0%)	
Host	Check	Cpu ti																															
srv-mysql-2	srv-mysql-2	189m																															
srv-web-1	Kernel Stats	180m																															
SERVER1_CUSTOMER2	Kernel Stats	179m																															
ht1	SSH Connexion	170m																															
SERVER1_CUSTOMER2	DiskStats	160m																															
0 - 50ms	50 - 100ms	100 - 200ms	200 - 300ms	300 - 400ms	+																												
182 (35%)	36 (6%)	301 (57%)	0 (0%)	0 (0%)																													

This check has a Result, which shows some information about the statistics of the Poller.


The long Result then provides a summary table with additional data.

Performance data

A check can also provide performance data. This is data returned by the check that will be stored and can be reused to draw graphs for example.

The "Poller - Performance" check returns data on CPU usage, CPU load, and the number of checks the CPU is able to perform.

This performance data is stored and can later be used to obtain a curve like the following one:



Status Confirmation

In order to confirm that the status returned when verifying a check is reliable, Shinken Enterprise can perform several checks of the same check..

- If a check returns an OK status, Shinken considers directly that this status is reliable.
- If, on the other hand, the verification of the check returns a status different from OK, Shinken restarts the check to confirm by a new verification that the status is different from OK.
 - The reason for this re-check is to prevent the sending of notifications on an uncertain state.
 - We then wait until we are assured that there is an incident before notifying users.
(more details on notifications in the associated page: [Basic - notifications](#)).