



This PDF is generated from authoritative online content, and is provided for convenience only. This PDF cannot be used for legal purposes. For authoritative understanding of what is and is not supported, always use the online content. To copy code samples, always use the online content.

Genesys Knowledge Center Deployment Guide

Load-Balancing Configuration

12/18/2025

Load-Balancing Configuration

Deploying a Cluster

Important

Whenever you deploy a Knowledge Center Server instance, you must configure a Knowledge Center Cluster, even if you only plan on having one server.

Knowledge Center Cluster stores all of the settings and data that are shared by each of the Knowledge Center Server instances that reside within it. This makes it pretty easy to add additional servers as your knowledge needs grow.

Knowledge Center Cluster also serves as the entry point to all client requests sent to Knowledge Center Servers. The cluster application in Genesys Administrator needs to be configured to point to the host and port of the load balancer that will distribute these requests among your Knowledge Center Servers.

Important

If you only have one server deployed in your cluster, you can configure the cluster application to point directly to the host and port of that server.

Configuring Your Load-Balancer Solution

Let's take a look at how you might configure your load balancer to distribute requests between servers. This sample uses an Apache load balancer.

Important

Genesys recommends that you use a round-robin approach to balancing for Knowledge Center Server. If want to use a cluster of Knowledge Center CMS instances you'll need to use a sticky session strategy for load-balancing in order to keep authorized user on the same node.

Important

If you need more information about load balancing in a Genesys environment, the Genesys Web Engagement [Load Balancing](#) page provides some useful background information.

Prerequisites

- Several Knowledge Center Servers/Knowledge Center CMS should be installed. These servers will be used as cluster nodes (node1, node2, node3, and so on)
- You must have a Genesys Administrator application type Application Cluster
- An application created for each Knowledge Center Servers/Knowledge Center CMS in cluster
- All Knowledge Center Server/Knowledge Center CMS applications should be connected to the application cluster

Start

1. Install Nginx HTTP and reverse proxy server (<https://nginx.org/en/>). The port and host of the installed load balancer should be used in the Application Cluster application in Genesys Administrator.
2. Configure nginx as the load-balancer for your cluster. See below for an example of configuring from a nginx.conf file.
 - Pay close attention to the upstream cluster_nodes section; here you can add a list of nodes and then select the balancing method.
 - For Knowledge Center Servers, we recommend using least_conn balancing method (http://nginx.org/en/docs/http/nginx_http_upstream_module.html#least_conn)
 - For Knowledge Center CMS, we recommend using ip_hash (http://nginx.org/en/docs/http/nginx_http_upstream_module.html#ip_hash)
3. Restart nginx, then restart all of your nodes.
4. All requests that are sent to nginx will be distributed to your cluster nodes.

```
user nginx;
worker_processes 1;

error_log /var/log/nginx/error.log warn;
pid /var/run/nginx.pid;

events {
    worker_connections 1024;
}

http {
    include /etc/nginx/mime.types;
    default_type application/octet-stream;

    log_format main '$remote_addr - $remote_user [$time_local] "$request" '
        '$status $body_bytes_sent "$http_referer" '
        '"$http_user_agent" "$http_x_forwarded_for"';

    access_log /var/log/nginx/access.log main;
```

```
sendfile on;
#tcp_nopush on;

keepalive_timeout 65;

#gzip on;

include /etc/nginx/conf.d/*.conf;

server {
    listen 9999;
    underscores_in_headers on;
    location / {
        proxy_pass http://cluster_nodes;
    }
}

upstream cluster_nodes {
    ip_hash;

    server gkc-node-1:9010;
    server gkc-node-1:9010;
}
}
```

End

Here are couple of sample requests:

- Request to a specific node: http://host_node_1:port_node_1/gks-sample-ui
- Request to the cluster, which will be distributed to any appropriate node:
http://host_load_balacer:port_load_balancer/gks-sample-ui