I. Nazwa serwera: app1

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	app1
IP / Domena	10.90.54.201
OS	CentOS 7
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	50 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	aplikacyjny (portal)

2. Konfiguracja .hosts

```
10.90.54.111 uwt1
10.90.54.112 uwt2

10.90.54.201 app1
10.90.54.202 app2
10.90.54.203 app3
10.90.54.204 app4
10.90.54.209 appcelery

10.90.54.211 storage1
10.90.54.212 storage2

10.90.54.221 db1

# short circuit
10.90.54.200 epodreczniki.pl www.epodreczniki.pl api.epodreczniki.pl preview.epodreczniki.pl static.epodreczniki.pl
epodreczniki.pcss.pl
10.90.54.200 repo.epodreczniki.pl content.epodreczniki.pl user.epodreczniki.pl
10.90.54.31 rt.epo.pl rt.epodreczniki.pl
10.90.54.250 avv.epodreczniki.pl
91.224.214.98 smtp.epodreczniki.pl
127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6
```

3. Interfejsy sieciowe

eth0 10.90.54.201

4. Zainstalowane usługi

a. Ngnix

Parametr	Wartość
Lokalizacja instalacji	/etc/nginx/
Zawartość pliku konfiguracyjnego #1	/etc/nginx/nginx.conf

```
# For more information on configuration, see:
#   * Official English Documentation: http://nginx.org/en/docs/
#   * Official Russian Documentation: http://nginx.org/ru/docs/

user nginx;
worker_processes auto;
error_log /var/log/nginx/error.log;
pid /run/nginx.pid;

# Load dynamic modules. See /usr/share/nginx/README.dynamic.
#include /usr/share/nginx/modules/*.conf;
load_module "modules/ngx_http_headers_more_filter_module.so";
events {
    worker_connections 1024;
}

http {
    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;
    tcp_nodelay         on;
    keepalive_timeout   65;
    types_hash_max_size 2048;

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

    # Load modular configuration files from the /etc/nginx/conf.d directory.
    # See http://nginx.org/en/docs/ngx_core_module.html#include
    # for more information.
    include /etc/nginx/conf.d/*.conf;

#    server {
#        listen       80 default_server;
#        listen       [::]:80 default_server;
#        server_name  _;
#        root         /usr/share/nginx/html;
#
#        # Load configuration files for the default server block.
#        include /etc/nginx/default.d/*.conf;
#
#        location / {
#        }
#
#        error_page 404 /404.html;
#            location = /40x.html {
#        }
#
#        error_page 500 502 503 504 /50x.html;
#            location = /50x.html {
#        }
#    }

# Settings for a TLS enabled server.
```

```
#
#   server {
#       listen       443 ssl http2 default_server;
#       listen       [::]:443 ssl http2 default_server;
#       server_name  _;
#       root         /usr/share/nginx/html;
#
#       ssl_certificate "/etc/pki/nginx/server.crt";
#       ssl_certificate_key "/etc/pki/nginx/private/server.key";
#       ssl_session_cache shared:SSL:1m;
#       ssl_session_timeout  10m;
#       ssl_ciphers HIGH:!aNULL:!MD5;
#       ssl_prefer_server_ciphers on;
#
#       # Load configuration files for the default server block.
#       include /etc/nginx/default.d/*.conf;
#
#       location / {
#       }
#
#       error_page 404 /404.html;
#           location = /40x.html {
#       }
#
#       error_page 500 502 503 504 /50x.html;
#           location = /50x.html {
#       }
#   }

}
```

Zawartość pliku konfiguracyjnego #2

/etc/nginx/conf.d

```
 upstream django {
   server 127.0.0.1:8000;
}

##
## epodreczniki.pl
##
server {
   listen 80;
   server_name epodreczniki.pl www.epodreczniki.pl;
....
   access_log /var/log/nginx/epodreczniki.pl_access.log;
   error_log  /var/log/nginx/epodreczniki.pl_error.log;
....
   location /portal/begin {
      root /opt/epodreczniki/epodreczniki-portal/local/media;
#       add_header 'Access-Control-Allow-Origin' *;
<------>more_set_headers 'Access-Control-Allow-Origin: *';
   }
....
   location / {
      more_set_headers 'Access-Control-Allow-Origin: *';
<------>client_body_timeout   300s;
      proxy_connect_timeout 300s;
      proxy_send_timeout    500s;
      proxy_set_header      Host www.epodreczniki.pl;
      #uwsgi_read_timeout    500;

      proxy_pass            http://django;
   }
}

##
## api.epodreczniki.pl
##
server {
   listen 80;
   server_name api.epodreczniki.pl www.api.epodreczniki.pl;
```

```
    access_log /var/log/nginx/api.epodreczniki.pl_access.log;
    error_log /var/log/nginx/api.epodreczniki.pl_error.log;

    location / {
 #      add_header       Access-Control-Allow-Origin *;
        more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_set_header Host api.epodreczniki.pl;
        proxy_pass       http://django;
    }
}



##
## edit.epodreczniki.pl
##
server {
    listen 80;
    server_name edit.epodreczniki.pl www.edit.epodreczniki.pl;
    access_log /var/log/nginx/edit.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/edit.epodreczniki.pl_error.log;

    location / {
 #      add_header       'Access-Control-Allow-Origin' *;
        proxy_set_header Host edit.epodreczniki.pl;
        proxy_pass       http://django;
    }
}


##
## preview.epodreczniki.pl
##
server {
    listen 80;
    server_name preview.epodreczniki.pl www.preview.epodreczniki.pl;

    access_log /var/log/nginx/preview.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/preview.epodreczniki.pl_error.log;

    proxy_send_timeout    300s;
    proxy_connect_timeout 300s;
    proxy_read_timeout    500s;

    location / {
        proxy_set_header Host preview.epodreczniki.pl;
<------>#more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_pass       http://django;
    }
}


##
## static.epodreczniki.pl
##
server {
    listen 80;
    server_name www.static.epodreczniki.pl static.epodreczniki.pl;

    access_log /var/log/nginx/static.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/static.epodreczniki.pl_error.log;

    location /portal {
  #     add_header 'Access-Control-Allow-Origin' *;
<------>more_set_headers 'Access-Control-Allow-Origin: *';
        alias      /opt/epodreczniki/epodreczniki-portal/portal/collected-static;
    }
.
    location /portal/other {
    #    add_header 'Access-Control-Allow-Origin' *;
<------>more_set_headers 'Access-Control-Allow-Origin: *';
        alias      /opt/epodreczniki/epodreczniki-portal/portal/static/other;
    }
```

```
    location / {
        #add_header       'Access-Control-Allow-Origin' *;
<------> more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_set_header Host static.epodreczniki.pl;
        proxy_pass       http://django;
    }
}
##
## content.epodreczniki.pl
##
server {
    listen 80;
    server_name content.epodreczniki.pl www.content.epodreczniki.pl content.epo.pl;

    access_log /var/log/nginx/content.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/content.epodreczniki.pl_error.log;

    location /content {
      alias      /opt/epodreczniki/epodreczniki-portal/portal/static/repository/content;
       #add_header 'Access-Control-Allow-Origin' '*';
<------>more_set_headers 'Access-Control-Allow-Origin: *';
      error_page 404 =  @err404;
      error_page 405 =  @err405;
   }

   location /global/libraries {
       alias      /opt/epodreczniki/epodreczniki-portal/portal/static/3rdparty;
       #add_header 'Access-Control-Allow-Origin' '*';
<------>more_set_headers 'Access-Control-Allow-Origin: *';
    }

    location / {
       #add_header       'Access-Control-Allow-Origin' '*';
<------>more_set_headers 'Access-Control-Allow-Origin: *';
       proxy_set_header Host content.epodreczniki.pl;
       proxy_pass       http://django;
    }

  location @err405 {
              #add_header   'Access-Control-Allow-Origin' '*';
       #more_set_headers 'Access-Control-Allow-Origin: *';
       proxy_set_header Host content.epodreczniki.pl;
       proxy_pass       http://django;
       #       uwsgi_pass django;
       #include uwsgi_params;
       #uwsgi_param     SCRIPT_NAME '';

  }
...
   location @err404 {
<------>     more_set_headers 'Access-Control-Allow-Origin: *';
          return  301 http://preview.epodreczniki.pl$request_uri;
       }..<--->
}

##
## user.epodreczniki.pl
##
server {
    listen 80;
    server_name user.epodreczniki.pl www.user.epodreczniki.pl;

    access_log /var/log/nginx/user.epodreczniki.pl_access.log;
    error_log /var/log/nginx/user.epodreczniki.pl_error.log;

    location / {
       proxy_set_header Host user.epodreczniki.pl;
       proxy_pass       http://django;
    }
}

##
```

```
## search.epodreczniki.pl
##
server {
   listen 80;
   server_name search.epodreczniki.pl www.search.epodreczniki.pl;

   access_log /var/log/nginx/search.epodreczniki.pl_access.log;
   error_log  /var/log/nginx/search.epodreczniki.pl_error.log;

   location / {
      proxy_set_header Host search.epodreczniki.pl;
      proxy_pass       http://django;
   }
}

##
## repo.epodreczniki.pl
##
server {
   listen 80;
   server_name repo.epodreczniki.pl www.repo.epodreczniki.pl;

   access_log /var/log/nginx/repo.epodreczniki.pl_access.log;
   error_log  /var/log/nginx/repo.epodreczniki.pl_error.log;

   location / {
         proxy_set_header Host repo.epodreczniki.pl;
<------><------>more_set_headers 'Access-Control-Allow-Origin: *';
         proxy_pass       http://django;
   }
}



##
## dev.epodreczniki.pl
##
#server {
#      listen 80;
#      server_name dev.epodreczniki.pl www.dev.epodreczniki.pl;
#      access_log /var/log/nginx/dev.epodreczniki.pl_access.log;
#      error_log /var/log/nginx/dev.epodreczniki.pl_error.log;
#       location / {
#             proxy_set_header Host dev.epodreczniki.pl;
#             proxy_pass http://127.0.0.1:8000;
#             add_header Access-Control-Allow-Origin '*';
#       }
#}
```

I. Nazwa serwera: app2

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	app2
IP / Domena	10.90.54.202
OS	CentOS 7
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	50 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	aplikacyjny (portal)

2. Konfiguracja .hosts

```
10.90.54.111 uwt1
10.90.54.112 uwt2

10.90.54.201 app1
10.90.54.202 app2
10.90.54.209 appcelery

10.90.54.211 storage1
10.90.54.212 storage2

10.90.54.221 db1

# short circuit
10.90.54.200 epodreczniki.pl www.epodreczniki.pl api.epodreczniki.pl preview.epodreczniki.pl static.epodreczniki.pl
10.90.54.200 repo.epodreczniki.pl content.epodreczniki.pl user.epodreczniki.pl
10.90.54.31 rt.epo.pl rt.epodreczniki.pl
10.90.54.250 avv.epodreczniki.pl

127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6
```

3. Interfejsy sieciowe

eth0 10.90.54.202

4. Zainstalowane usługi

a. Nginx

Parametr	Wartość
Lokalizacja instalacji	/etc/nginx/
Zawartość pliku konfiguracyjnego #1	/etc/nginx/nginx.conf

```
# For more information on configuration, see:
#   * Official English Documentation: http://nginx.org/en/docs/
#   * Official Russian Documentation: http://nginx.org/ru/docs/

user nginx;
worker_processes auto;
error_log /var/log/nginx/error.log;
pid /run/nginx.pid;

# Load dynamic modules. See /usr/share/nginx/README.dynamic.
load_module "modules/ngx_http_headers_more_filter_module.so";
#include /usr/share/nginx/modules/*.conf;

events {
   worker_connections 1024;
}

http {
   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;
   tcp_nodelay         on;
   keepalive_timeout   65;
   types_hash_max_size 2048;

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

   # Load modular configuration files from the /etc/nginx/conf.d directory.
   # See http://nginx.org/en/docs/ngx_core_module.html#include
   # for more information.
   include /etc/nginx/conf.d/*.conf;

#   server {
#       listen       80 default_server;
#       listen       [::]:80 default_server;
#       server_name  _;
#       root         /usr/share/nginx/html;
#
#       # Load configuration files for the default server block.
#       include /etc/nginx/default.d/*.conf;
#
#       location / {
#       }
#
#       error_page 404 /404.html;
#           location = /40x.html {
#       }
#
#       error_page 500 502 503 504 /50x.html;
#           location = /50x.html {
#       }
#   }
```

Zawartość pliku konfiguracyjnego #2

/etc/nginx/conf.d/epodreczniki.conf

```
upstream django {
   server 127.0.0.1:8000;
}

##
## epodreczniki.pl
##
server {
   listen 80;
   server_name epodreczniki.pl www.epodreczniki.pl;
....
   access_log /var/log/nginx/epodreczniki.pl_access.log;
   error_log  /var/log/nginx/epodreczniki.pl_error.log;
....
   location /portal/begin {
      root /opt/epodreczniki/epodreczniki-portal/local/media;
#       add_header 'Access-Control-Allow-Origin' *;
<------>more_set_headers 'Access-Control-Allow-Origin: *';
   }
....
   location / {
      more_set_headers 'Access-Control-Allow-Origin: *';
<------>client_body_timeout   300s;
      proxy_connect_timeout 300s;
      proxy_send_timeout    500s;
      proxy_set_header      Host www.epodreczniki.pl;
      #uwsgi_read_timeout    500;

      proxy_pass            http://django;
   }
}

##
## api.epodreczniki.pl
##
server {
   listen 80;
   server_name api.epodreczniki.pl www.api.epodreczniki.pl;

   access_log /var/log/nginx/api.epodreczniki.pl_access.log;
   error_log /var/log/nginx/api.epodreczniki.pl_error.log;

   location / {
#       add_header       Access-Control-Allow-Origin *;
      more_set_headers 'Access-Control-Allow-Origin: *';
      proxy_set_header Host api.epodreczniki.pl;
      proxy_pass       http://django;
   }
}



##
## edit.epodreczniki.pl
##
server {
   listen 80;
   server_name edit.epodreczniki.pl www.edit.epodreczniki.pl;
   access_log /var/log/nginx/edit.epodreczniki.pl_access.log;
   error_log  /var/log/nginx/edit.epodreczniki.pl_error.log;

   location / {
#       add_header       'Access-Control-Allow-Origin' *;
      proxy_set_header Host edit.epodreczniki.pl;
      proxy_pass       http://django;
   }
}


##
```

```
## preview.epodreczniki.pl
##
server {
   listen 80;
   server_name preview.epodreczniki.pl www.preview.epodreczniki.pl;

   access_log /var/log/nginx/preview.epodreczniki.pl_access.log;
   error_log  /var/log/nginx/preview.epodreczniki.pl_error.log;

   proxy_send_timeout    300s;
   proxy_connect_timeout 300s;
   proxy_read_timeout    500s;

   location / {
      proxy_set_header Host preview.epodreczniki.pl;
<------>#more_set_headers 'Access-Control-Allow-Origin: *';
      proxy_pass       http://django;
   }
}


##
## static.epodreczniki.pl
##
server {
   listen 80;
   server_name www.static.epodreczniki.pl static.epodreczniki.pl;

   access_log /var/log/nginx/static.epodreczniki.pl_access.log;
   error_log  /var/log/nginx/static.epodreczniki.pl_error.log;

   location /portal {
  #     add_header 'Access-Control-Allow-Origin' *;
<------>more_set_headers 'Access-Control-Allow-Origin: *';
      alias      /opt/epodreczniki/epodreczniki-portal/portal/collected-static;
   }
.
   location /portal/other {
   #    add_header 'Access-Control-Allow-Origin' *;
<------>more_set_headers 'Access-Control-Allow-Origin: *';
      alias      /opt/epodreczniki/epodreczniki-portal/portal/static/other;
   }

   location / {
      #add_header       'Access-Control-Allow-Origin' *;
<------> more_set_headers 'Access-Control-Allow-Origin: *';
      proxy_set_header Host static.epodreczniki.pl;
      proxy_pass       http://django;
   }
}
##
## content.epodreczniki.pl
##
server {
   listen 80;
   server_name content.epodreczniki.pl www.content.epodreczniki.pl content.epo.pl;

   access_log /var/log/nginx/content.epodreczniki.pl_access.log;
   error_log  /var/log/nginx/content.epodreczniki.pl_error.log;

   location /content {
      alias      /opt/epodreczniki/epodreczniki-portal/portal/static/repository/content;
      #add_header 'Access-Control-Allow-Origin' '*';
<------>more_set_headers 'Access-Control-Allow-Origin: *';
     error_page 404 =  @err404;
     error_page 405 =  @err405;
   }

   location /global/libraries {
      alias      /opt/epodreczniki/epodreczniki-portal/portal/static/3rdparty;
      #add_header 'Access-Control-Allow-Origin' '*';
<------>more_set_headers 'Access-Control-Allow-Origin: *';
   }
```

```
    location / {
        #add_header       'Access-Control-Allow-Origin' '*';
<------>more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_set_header Host content.epodreczniki.pl;
        proxy_pass       http://django;
    }
  location @err405 {
              #add_header      'Access-Control-Allow-Origin' '*';
        #more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_set_header Host content.epodreczniki.pl;
        proxy_pass       http://django;
        #       uwsgi_pass django;
        #include uwsgi_params;
        #uwsgi_param     SCRIPT_NAME '';

  }...
    location @err404 {
<------>     more_set_headers 'Access-Control-Allow-Origin: *';
          return  301 http://preview.epodreczniki.pl$request_uri;
        }..<--->
}

##
## user.epodreczniki.pl
##
server {
    listen 80;
    server_name user.epodreczniki.pl www.user.epodreczniki.pl;

    access_log /var/log/nginx/user.epodreczniki.pl_access.log;
    error_log /var/log/nginx/user.epodreczniki.pl_error.log;

    location / {
        proxy_set_header Host user.epodreczniki.pl;
        #more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_pass       http://django;
    }
}

##
## search.epodreczniki.pl
##
server {
    listen 80;
    server_name search.epodreczniki.pl www.search.epodreczniki.pl;

    access_log /var/log/nginx/search.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/search.epodreczniki.pl_error.log;

    location / {
        proxy_set_header Host search.epodreczniki.pl;
        proxy_pass       http://django;
    }
}

##
## repo.epodreczniki.pl
##
server {
    listen 80;
    server_name repo.epodreczniki.pl www.repo.epodreczniki.pl;

    access_log /var/log/nginx/repo.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/repo.epodreczniki.pl_error.log;

    location / {
          proxy_set_header Host repo.epodreczniki.pl;
          #add_header       'Access-Control-Allow-Origin' '*';
<------><------>more_set_headers 'Access-Control-Allow-Origin: *';
          proxy_pass       http://django;
    }
}
```

I. Nazwa serwera: app3

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	app3
IP / Domena	10.90.54.203
OS	CentOS 7
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	50 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	aplikacyjny (portal)

2. Konfiguracja .hosts

```
10.90.54.111 uwt1
10.90.54.112 uwt2

10.90.54.201 app1
10.90.54.202 app2
10.90.54.203 app3
10.90.54.204 app4
10.90.54.209 appcelery

10.90.54.211 storage1
10.90.54.212 storage2

10.90.54.221 db1

# short circuit
10.90.54.200 epodreczniki.pl www.epodreczniki.pl api.epodreczniki.pl preview.epodreczniki.pl static.epodreczniki.pl
10.90.54.200 repo.epodreczniki.pl content.epodreczniki.pl user.epodreczniki.pl
10.90.54.31 rt.epo.pl rt.epodreczniki.pl
10.90.54.250 avv.epodreczniki.pl

127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6
```

3. Interfejsy sieciowe

eth0 10.90.54.203

4. Zainstalowane usługi

   a. Ngnix

Parametr	Wartość
Lokalizacja instalacji	/etc/nginx/

Zawartość pliku konfiguracyjnego #1	/etc/nginx/nginx.conf
	
	# For more information on configuration, see:
	# * Official English Documentation: http://nginx.org/en/docs/
	# * Official Russian Documentation: http://nginx.org/ru/docs/
	
	user nginx;
	worker_processes auto;
	error_log /var/log/nginx/error.log;
	pid /run/nginx.pid;
	
	# Load dynamic modules. See /usr/share/nginx/README.dynamic.
	load_module "modules/ngx_http_headers_more_filter_module.so";
	events {
	  worker_connections 1024;
	}
	
	http {
	  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;
	  tcp_nodelay on;
	  keepalive_timeout 65;
	  types_hash_max_size 2048;
	
	  include /etc/nginx/mime.types;
	  default_type application/octet-stream;
	
	  include /etc/nginx/conf.d/*.conf;
	}

Zawartość pliku konfiguracyjnego #2	/etc/nginx/conf.d/epodreczniki.conf
	
	upstream django {
	  server 127.0.0.1:8000;
	}
	
	##
	## epodreczniki.pl
	##
	server {
	  listen 80;
	  server_name epodreczniki.pl www.epodreczniki.pl;
	....
	  access_log /var/log/nginx/epodreczniki.pl_access.log;
	  error_log /var/log/nginx/epodreczniki.pl_error.log;
	....
	  location /portal/begin {
	  root /opt/epodreczniki/epodreczniki-portal/local/media;
	# add_header 'Access-Control-Allow-Origin' *;
	<------>more_set_headers 'Access-Control-Allow-Origin: *' always;
	  }
	....
	  location / {
	  more_set_headers 'Access-Control-Allow-Origin: *' always;
	<------>client_body_timeout 300s;
	  proxy_connect_timeout 300s;
	  proxy_send_timeout 500s;

```
        proxy_set_header      Host www.epodreczniki.pl;
        #uwsgi_read_timeout    500;

        proxy_pass            http://django;
    }
}

##
## api.epodreczniki.pl
##
server {
    listen 80;
    server_name api.epodreczniki.pl www.api.epodreczniki.pl;

    access_log /var/log/nginx/api.epodreczniki.pl_access.log;
    error_log /var/log/nginx/api.epodreczniki.pl_error.log;

    location / {
 #       add_header       Access-Control-Allow-Origin *;
        more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_set_header Host api.epodreczniki.pl;
        proxy_pass       http://django;
    }
}



##
## edit.epodreczniki.pl
##
server {
    listen 80;
    server_name edit.epodreczniki.pl www.edit.epodreczniki.pl;
    access_log /var/log/nginx/edit.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/edit.epodreczniki.pl_error.log;

    location / {
 #      add_header       'Access-Control-Allow-Origin' *;
        proxy_set_header Host edit.epodreczniki.pl;
        proxy_pass       http://django;
    }
}


##
## preview.epodreczniki.pl
##
server {
    listen 80;
    server_name preview.epodreczniki.pl www.preview.epodreczniki.pl;

    access_log /var/log/nginx/preview.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/preview.epodreczniki.pl_error.log;

    proxy_send_timeout    300s;
    proxy_connect_timeout 300s;
    proxy_read_timeout    500s;

    location / {
        proxy_set_header Host preview.epodreczniki.pl;
<------>#more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_pass       http://django;
    }
}


##
## static.epodreczniki.pl
##
server {
    listen 80;
    server_name www.static.epodreczniki.pl static.epodreczniki.pl;
```

```
    access_log /var/log/nginx/static.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/static.epodreczniki.pl_error.log;

    location /portal {
   #     add_header 'Access-Control-Allow-Origin' *;
<------>more_set_headers 'Access-Control-Allow-Origin: *';
        alias     /opt/epodreczniki/epodreczniki-portal/portal/collected-static;
    }
.
    location /portal/other {
    #    add_header 'Access-Control-Allow-Origin' *;
<------>more_set_headers 'Access-Control-Allow-Origin: *';
        alias     /opt/epodreczniki/epodreczniki-portal/portal/static/other;
    }

    location / {
        #add_header      'Access-Control-Allow-Origin' *;
<------> more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_set_header Host static.epodreczniki.pl;
        proxy_pass      http://django;
    }
}

##
## content.epodreczniki.pl
##
server {
    listen 80;
    server_name content.epodreczniki.pl www.content.epodreczniki.pl content.epo.pl;

    access_log /var/log/nginx/content.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/content.epodreczniki.pl_error.log;

    location /content {
        alias     /opt/epodreczniki/epodreczniki-portal/portal/static/repository/content;
        #add_header 'Access-Control-Allow-Origin' '*';
<------>more_set_headers 'Access-Control-Allow-Origin: *';
      error_page 404 =  @err404;
     error_page 405 =  @err405;
    }

   location /global/libraries {
        alias     /opt/epodreczniki/epodreczniki-portal/portal/static/3rdparty;
        #add_header 'Access-Control-Allow-Origin' '*';
<------>more_set_headers 'Access-Control-Allow-Origin: *';
    }

    location / {
        #add_header      'Access-Control-Allow-Origin' '*';
<------>more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_set_header Host content.epodreczniki.pl;
        proxy_pass      http://django;
    }

  location @err405 {
              #add_header   'Access-Control-Allow-Origin' '*';
        #more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_set_header Host content.epodreczniki.pl;
        proxy_pass      http://django;
        #       uwsgi_pass django;
        #include uwsgi_params;
        #uwsgi_param     SCRIPT_NAME '';

  }...
   location @err404 {
<------>    more_set_headers 'Access-Control-Allow-Origin: *';
          return  301 http://preview.epodreczniki.pl$request_uri;
          #return 301 http://preview.epodreczniki.pl$uri$is_args$args;
       }..<--->
}

##
## user.epodreczniki.pl
```

```
##
server {
   listen 80;
   server_name user.epodreczniki.pl www.user.epodreczniki.pl;

   access_log /var/log/nginx/user.epodreczniki.pl_access.log;
   error_log /var/log/nginx/user.epodreczniki.pl_error.log;

   location / {
      #more_set_headers 'Access-Control-Allow-Origin: *';
      proxy_set_header Host user.epodreczniki.pl;
      proxy_pass       http://django;
.....
   }
}

##
## search.epodreczniki.pl
##
server {
   listen 80;
   server_name search.epodreczniki.pl www.search.epodreczniki.pl;

   access_log /var/log/nginx/search.epodreczniki.pl_access.log;
   error_log  /var/log/nginx/search.epodreczniki.pl_error.log;

   location / {
      proxy_set_header Host search.epodreczniki.pl;
      proxy_pass       http://django;
   }
}

##
## repo.epodreczniki.pl
##
server {
   listen 80;
   server_name repo.epodreczniki.pl www.repo.epodreczniki.pl;

   access_log /var/log/nginx/repo.epodreczniki.pl_access.log;
   error_log  /var/log/nginx/repo.epodreczniki.pl_error.log;

   location / {
         proxy_set_header Host repo.epodreczniki.pl;
         #add_header       'Access-Control-Allow-Origin' '*';
<------><------>more_set_headers 'Access-Control-Allow-Origin: *';
         proxy_pass       http://django;
   }
}



##
## dev.epodreczniki.pl
##
#server {
#       listen 80;
#       server_name dev.epodreczniki.pl www.dev.epodreczniki.pl;
#       access_log /var/log/nginx/dev.epodreczniki.pl_access.log;
#       error_log /var/log/nginx/dev.epodreczniki.pl_error.log;
#        location / {
#              proxy_set_header Host dev.epodreczniki.pl;
#              proxy_pass http://127.0.0.1:8000;
#              add_header Access-Control-Allow-Origin '*';
#       }
#}
```

I. Nazwa serwera: app4

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	app4
IP / Domena	10.90.54.204
OS	CentOS 7
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	50 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	aplikacyjny (portal)

2. Konfiguracja .hosts

```
10.90.54.111 uwt1
10.90.54.112 uwt2

10.90.54.201 app1
10.90.54.202 app2
10.90.54.203 app3
10.90.54.204 app4
10.90.54.209 appcelery

10.90.54.211 storage1
10.90.54.212 storage2

10.90.54.221 db1

# short circuit
10.90.54.200 epodreczniki.pl www.epodreczniki.pl api.epodreczniki.pl preview.epodreczniki.pl static.epodreczniki.pl
10.90.54.200 repo.epodreczniki.pl content.epodreczniki.pl user.epodreczniki.pl
10.90.54.31 rt.epo.pl rt.epodreczniki.pl
10.90.54.250 avv.epodreczniki.pl

127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6
```

3. Interfejsy sieciowe

eth0 10.90.54.204

4. Zainstalowane usługi

a. Ngnix

Parametr	Wartość
Lokalizacja instalacji	/etc/nginx/

Zawartość pliku konfiguracyjnego #1	/etc/nginx/nginx.conf
	
	load_module "modules/ngx_http_headers_more_filter_module.so";
	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;
	}

Zawartość pliku konfiguracyjnego #2	/etc/nginx/conf.d/epodreczniki.conf
	
	upstream django {
	# fail_timeout=0 means we always retry an upstream even if it failed
	# to return a good HTTP response (in case the Unicorn master nukes a
	# single worker for timing out).
	
	server unix:///tmp/epo-portal.sock fail_timeout=0;
	}
	
	##
	## epodreczniki.pl
	##
	server {
	listen 80;
	server_name epodreczniki.pl www.epodreczniki.pl;
	....
	access_log /var/log/nginx/epodreczniki.pl_access.log;
	error_log /var/log/nginx/epodreczniki.pl_error.log;
	....
	location /portal/begin {
	root /opt/epodreczniki/epodreczniki-portal/local/media;
	# add_header 'Access-Control-Allow-Origin' *;
	<------>more_set_headers 'Access-Control-Allow-Origin: *';
	}
	....
	location / {

```
        more_set_headers 'Access-Control-Allow-Origin: *';
<------>client_body_timeout   300s;
        proxy_connect_timeout 300s;
        proxy_send_timeout    500s;
        proxy_set_header      Host www.epodreczniki.pl;
        #uwsgi_read_timeout    500;

        #proxy_pass           http://django;
<------>uwsgi_pass django;
        include uwsgi_params;
        uwsgi_param     SCRIPT_NAME '';
    }
}

##
## api.epodreczniki.pl
##
server {
    listen 80;
    server_name api.epodreczniki.pl www.api.epodreczniki.pl;

    access_log /var/log/nginx/api.epodreczniki.pl_access.log;
    error_log /var/log/nginx/api.epodreczniki.pl_error.log;

    location / {
 #      add_header       Access-Control-Allow-Origin *;
        more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_set_header Host api.epodreczniki.pl;
        #proxy_pass       http://django;
.
       uwsgi_pass django;
        include uwsgi_params;
        uwsgi_param     SCRIPT_NAME '';

    }
}



##
## edit.epodreczniki.pl
##
server {
    listen 80;
    server_name edit.epodreczniki.pl www.edit.epodreczniki.pl;
    access_log /var/log/nginx/edit.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/edit.epodreczniki.pl_error.log;

    location / {
 #      add_header       'Access-Control-Allow-Origin' *;
        proxy_set_header Host edit.epodreczniki.pl;
        #proxy_pass       http://django;
<------>        uwsgi_pass django;
        include uwsgi_params;
        uwsgi_param     SCRIPT_NAME '';

    }
}


##
## preview.epodreczniki.pl
##
server {
    listen 80;
    server_name preview.epodreczniki.pl www.preview.epodreczniki.pl;

    access_log /var/log/nginx/preview.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/preview.epodreczniki.pl_error.log;

    proxy_send_timeout    300s;
    proxy_connect_timeout 300s;
    proxy_read_timeout    500s;
```

```
    location / {
        proxy_set_header Host preview.epodreczniki.pl;
<------>#more_set_headers 'Access-Control-Allow-Origin: *';
        #proxy_pass       http://django;
<------>        uwsgi_pass django;
        include uwsgi_params;
        uwsgi_param     SCRIPT_NAME '';

    }
}


##
## static.epodreczniki.pl
##
server {
    listen 80;
    server_name www.static.epodreczniki.pl static.epodreczniki.pl;

    access_log /var/log/nginx/static.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/static.epodreczniki.pl_error.log;

    location /portal {
  #     add_header 'Access-Control-Allow-Origin' *;
<------>more_set_headers 'Access-Control-Allow-Origin: *';
        alias      /opt/epodreczniki/epodreczniki-portal/portal/collected-static1;
    }
.
    location /portal/other {
    #    add_header 'Access-Control-Allow-Origin' *;
<------>more_set_headers 'Access-Control-Allow-Origin: *';
        alias      /opt/epodreczniki/epodreczniki-portal/portal/static1/other;
    }

    location / {
        #add_header       'Access-Control-Allow-Origin' *;
<------> more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_set_header Host static.epodreczniki.pl;
        #proxy_pass       http://django;
<------>        uwsgi_pass django;
        include uwsgi_params;
        uwsgi_param     SCRIPT_NAME '';

    }
}

##
## content.epodreczniki.pl
##
server {
    listen 80;
    server_name content.epodreczniki.pl www.content.epodreczniki.pl content.epo.pl;

    access_log /var/log/nginx/content.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/content.epodreczniki.pl_error.log;

    location /content {
        alias<->   /opt/epodreczniki/epodreczniki-portal/portal/static/repository/content;
        #add_header 'Access-Control-Allow-Origin' '*';
        more_set_headers 'Access-Control-Allow-Origin: *';
#       error_page 404 =  @err404;
        error_page 405 =  @err405;
    }

    location /global/libraries {
        alias<->   /opt/epodreczniki/epodreczniki-portal/portal/static/3rdparty;
        #add_header 'Access-Control-Allow-Origin' '*';
        more_set_headers 'Access-Control-Allow-Origin: *';
    }


    location / {
```

```
        #add_header       'Access-Control-Allow-Origin' '*';
<------>#more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_set_header Host content.epodreczniki.pl;
        #proxy_pass       http://django;
<------>        uwsgi_pass django;
        include uwsgi_params;
        uwsgi_param     SCRIPT_NAME '';

    }
..
  location @err405 {
              #add_header <->  'Access-Control-Allow-Origin' '*';
        #more_set_headers 'Access-Control-Allow-Origin: *';
        proxy_set_header Host content.epodreczniki.pl;
        #proxy_pass<--->   http://django;
             uwsgi_pass django;
        include uwsgi_params;
        uwsgi_param     SCRIPT_NAME '';

  }
.
   location @err404 {
<------>     more_set_headers 'Access-Control-Allow-Origin: *';
          return  301 http://preview.epodreczniki.pl$request_uri;
    #       return 301 http://preview.epodreczniki.pl$uri$is_args$args;
  }
}

##
## user.epodreczniki.pl
##
server {
    listen 80;
    server_name user.epodreczniki.pl www.user.epodreczniki.pl;

    access_log /var/log/nginx/user.epodreczniki.pl_access.log;
    error_log /var/log/nginx/user.epodreczniki.pl_error.log;

    location / {
        proxy_set_header Host user.epodreczniki.pl;
        #proxy_pass       http://django;
<------>        uwsgi_pass django;
        include uwsgi_params;
        uwsgi_param     SCRIPT_NAME '';

    }
}

##
## search.epodreczniki.pl
##
server {
    listen 80;
    server_name search.epodreczniki.pl www.search.epodreczniki.pl;

    access_log /var/log/nginx/search.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/search.epodreczniki.pl_error.log;

    location / {
        proxy_set_header Host search.epodreczniki.pl;
        #proxy_pass       http://django;

        uwsgi_pass django;
        include uwsgi_params;
        uwsgi_param     SCRIPT_NAME '';

    }
}

##
## repo.epodreczniki.pl
##
server {
```

```
    listen 80;
    server_name repo.epodreczniki.pl www.repo.epodreczniki.pl;

    access_log /var/log/nginx/repo.epodreczniki.pl_access.log;
    error_log  /var/log/nginx/repo.epodreczniki.pl_error.log;

    location / {
          proxy_set_header Host repo.epodreczniki.pl;
          #add_header      'Access-Control-Allow-Origin' '*';
<------><------>more_set_headers 'Access-Control-Allow-Origin: *';
          #proxy_pass      http://django;
<------>          uwsgi_pass django;
       include uwsgi_params;
       uwsgi_param     SCRIPT_NAME '';

    }
}



##
## dev.epodreczniki.pl
##
#server {
#       listen 80;
#       server_name dev.epodreczniki.pl www.dev.epodreczniki.pl;
#       access_log /var/log/nginx/dev.epodreczniki.pl_access.log;
#       error_log /var/log/nginx/dev.epodreczniki.pl_error.log;
#        location / {
#              proxy_set_header Host dev.epodreczniki.pl;
#              proxy_pass http://127.0.0.1:8000;
#              add_header Access-Control-Allow-Origin '*';
#       }
#}
```

I. Nazwa serwera: app-celery

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	app-celery
IP / Domena	10.90.54.209
OS	CentOS 7
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	50
Zewnętrzne IP:	NIE
Funkcja serwera:	Kolejka zadań

2. Konfiguracja .hosts

```
10.90.54.111 uwt1
10.90.54.112 uwt2

10.90.54.201 app1
10.90.54.202 app2
10.90.54.209 appcelery

10.90.54.211 storage1
10.90.54.212 storage2

10.90.54.221 db1

# internal proxy
10.90.54.200 epodreczniki.pl www.epodreczniki.pl api.epodreczniki.pl preview.epodreczniki.pl static.epodreczniki.pl
10.90.54.200 repo.epodreczniki.pl content.epodreczniki.pl user.epodreczniki.pl
10.90.54.31 rt.epo.pl rt.epodreczniki.pl

127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6
```

3. Interfejsy sieciowe

eth0 10.90.54.209

4. Zainstalowane usługi

a. Celery

Parametr	Wartość
Lokalizacja instalacji	/opt/epodreczniki/epodreczniki-portal/portal/portal/settings/deductions

Zawartość pliku konfiguracyjnego #1	import kombu
	
	BROKER_URL = None
	
	if not SURROUND_RUNNING_ON_PLATFORM:
	  if EPO_DEV_CELERY_MODE == 'eager':
	  BROKER_URL = 'django://'
	  CELERY_ALWAYS_EAGER = True
	  elif EPO_DEV_CELERY_MODE == 'django':
	  raise Exception('this is broken')
	  # RuntimeError: Conflicting 'queue' models in application 'kombu_transport_django': <class 'kombu.transport.django.models.Queue'> and <class 'djcelery.transport.models.Queue'>
	  BROKER_URL = 'django://'
	  INSTALLED_APPS += ('djcelery.transport',)
	  # INSTALLED_APPS += ('kombu.transport.django',)
	  elif EPO_DEV_CELERY_MODE == 'amqp':
	  # the catch call below works
	  pass
	  else:
	  raise Exception('invalid value of EPO_DEV_CELERY_MODE setting: %s' % EPO_DEV_CELERY_MODE)
	
	if BROKER_URL is None:
	  BROKER_URL = 'amqp://epo:epo@10.90.54.91:5672/epo'
	  BROKER_API = 'http://epo:epo@127.0.0.1:15672/api'
	
	CELERY_QUEUES = []
	
	for role in ["static", "app"]:
	  for host in HOSTS[role]:
	  queue = '%s-%s' % (role, host)
	  CELERY_QUEUES.append(kombu.Queue(queue, routing_key=queue))
	
	CELERY_QUEUES.append(kombu.Queue("static", routing_key='static'))
	CELERY_QUEUES.append(kombu.Queue("app", routing_key='app'))
	CELERY_QUEUES.append(kombu.Queue("publisher", routing_key='publisher'))

	CELERY_QUEUES.append(kombu.Queue("indexer", routing_key='indexer')) CELERY_QUEUES.append(kombu.Queue("mailer", routing_key='mailer')) CELERY_QUEUES.append(kombu.Queue("warmer", routing_key='warmer')) CELERY_QUEUES.append(kombu.Queue("purger", routing_key='purger')) CELERY_QUEUES.append(kombu.Queue("celery", routing_key='default')) CELERY_ACCEPT_CONTENT = ['pickle', 'json', 'msgpack', 'yaml']
Zawartość pliku konfiguracyjnego #2	<wklej treść pliku>

b. uwsgi

Parametr	Wartość
Lokalizacja instalacji	/opt/epodreczniki/uwsgi/epo-portal.ini

Zawartość pliku konfiguracyjnego #1	[uwsgi]
	chdir = /opt/epodreczniki/epodreczniki-portal/portal
	home = /opt/epodreczniki/venv/
	wsgi-file = wsgi.py
	
	#virtualenv = /opt/epodreczniki/venv
	#env = DJANGO_SETTINGS_MODULE=portal.settings.instances.simple
	
	buffer-size = 387608
	#listen = 2048
	#max-requests = 350000
	
	http = 127.0.0.1:8000
	
	# UWSGI workers config
	master = true
	die-on-term = true
	vacuum = true
	workers = 8
	enable-threads = true

Zawartość pliku konfiguracyjnego #2	<wklej treść pliku>

c. Aplikacja www

Parametr	Wartość
Lokalizacja instalacji	/etc/redis.conf
Zawartość pliku konfiguracyjnego #1	# Redis configuration file example. # # Note that in order to read the configuration file, Redis must be # started with the file path as first argument: # # ./redis-server /path/to/redis.conf # Note on units: when memory size is needed, it is possible to specify # it in the usual form of 1k 5GB 4M and so forth: # # 1k => 1000 bytes # 1kb => 1024 bytes # 1m => 1000000 bytes # 1mb => 1024*1024 bytes # 1g => 1000000000 bytes # 1gb => 1024*1024*1024 bytes # # units are case insensitive so 1GB 1Gb 1gB are all the same. ################################## INCLUDES ################################### # Include one or more other config files here. This is useful if you # have a standard template that goes to all Redis servers but also need # to customize a few per-server settings. Include files can include # other files, so use this wisely. # # Notice option "include" won't be rewritten by command "CONFIG REWRITE" # from admin or Redis Sentinel. Since Redis always uses the last processed # line as value of a configuration directive, you'd better put includes # at the beginning of this file to avoid overwriting config change at runtime. # # If instead you are interested in using includes to override configuration # options, it is better to use include as the last line. # # include /path/to/local.conf

```
# include /path/to/other.conf

################################## NETWORK
#####################################

# By default, if no "bind" configuration directive is specified, Redis
listens
# for connections from all the network interfaces available on the server.
# It is possible to listen to just one or multiple selected interfaces using
# the "bind" configuration directive, followed by one or more IP
addresses.
#
# Examples:
#
# bind 192.168.1.100 10.0.0.1
# bind 127.0.0.1 ::1
#
# ~~~ WARNING ~~~ If the computer running Redis is directly exposed
to the
# internet, binding to all the interfaces is dangerous and will expose the
# instance to everybody on the internet. So by default we uncomment the
# following bind directive, that will force Redis to listen only into
# the IPv4 lookback interface address (this means Redis will be able to
# accept connections only from clients running into the same computer it
# is running).
#
# IF YOU ARE SURE YOU WANT YOUR INSTANCE TO LISTEN
TO ALL THE INTERFACES
# JUST COMMENT THE FOLLOWING LINE.
#
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
~~~~~~~~~~~~~~~~~~~~
bind 127.0.0.1
bind 10.90.54.209

# Protected mode is a layer of security protection, in order to avoid that
# Redis instances left open on the internet are accessed and exploited.
#
# When protected mode is on and if:
#
# 1) The server is not binding explicitly to a set of addresses using the
#    "bind" directive.
# 2) No password is configured.
#
# The server only accepts connections from clients connecting from the
# IPv4 and IPv6 loopback addresses 127.0.0.1 and ::1, and from Unix
domain
# sockets.
```

```
#
# By default protected mode is enabled. You should disable it only if
# you are sure you want clients from other hosts to connect to Redis
# even if no authentication is configured, nor a specific set of interfaces
# are explicitly listed using the "bind" directive.
protected-mode yes

# Accept connections on the specified port, default is 6379 (IANA
#815344).
# If port 0 is specified Redis will not listen on a TCP socket.
port 6379

# TCP listen() backlog.
#
# In high requests-per-second environments you need an high backlog in
order
# to avoid slow clients connections issues. Note that the Linux kernel
# will silently truncate it to the value of /proc/sys/net/core/somaxconn so
# make sure to raise both the value of somaxconn and
tcp_max_syn_backlog
# in order to get the desired effect.
tcp-backlog 511

# Unix socket.
#
# Specify the path for the Unix socket that will be used to listen for
# incoming connections. There is no default, so Redis will not listen
# on a unix socket when not specified.
#
# unixsocket /tmp/redis.sock
# unixsocketperm 700

# Close the connection after a client is idle for N seconds (0 to disable)
timeout 0

# TCP keepalive.
#
# If non-zero, use SO_KEEPALIVE to send TCP ACKs to clients in
absence
# of communication. This is useful for two reasons:
#
# 1) Detect dead peers.
# 2) Take the connection alive from the point of view of network
#    equipment in the middle.
#
# On Linux, the specified value (in seconds) is the period used to send
ACKs.
# Note that to close the connection the double of the time is needed.
```

```
# On other kernels the period depends on the kernel configuration.
#
# A reasonable value for this option is 300 seconds, which is the new
# Redis default starting with Redis 3.2.1.
tcp-keepalive 300

################################## GENERAL
#####################################

# By default Redis does not run as a daemon. Use 'yes' if you need it.
# Note that Redis will write a pid file in /var/run/redis.pid when
daemonized.
daemonize no

# If you run Redis from upstart or systemd, Redis can interact with your
# supervision tree. Options:
#   supervised no      - no supervision interaction
#   supervised upstart - signal upstart by putting Redis into SIGSTOP
mode
#   supervised systemd - signal systemd by writing READY=1 to
$NOTIFY_SOCKET
#   supervised auto    - detect upstart or systemd method based on
#                        UPSTART_JOB or NOTIFY_SOCKET environment
variables
# Note: these supervision methods only signal "process is ready."
#       They do not enable continuous liveness pings back to your
supervisor.
supervised no

# If a pid file is specified, Redis writes it where specified at startup
# and removes it at exit.
#
# When the server runs non daemonized, no pid file is created if none is
# specified in the configuration. When the server is daemonized, the pid
file
# is used even if not specified, defaulting to "/var/run/redis.pid".
#
# Creating a pid file is best effort: if Redis is not able to create it
# nothing bad happens, the server will start and run normally.
pidfile /var/run/redis/redis.pid

# Specify the server verbosity level.
# This can be one of:
# debug (a lot of information, useful for development/testing)
# verbose (many rarely useful info, but not a mess like the debug level)
# notice (moderately verbose, what you want in production probably)
# warning (only very important / critical messages are logged)
loglevel notice
```

```
# Specify the log file name. Also the empty string can be used to force
# Redis to log on the standard output. Note that if you use standard
# output for logging but daemonize, logs will be sent to /dev/null
logfile /var/log/redis/redis.log

# To enable logging to the system logger, just set 'syslog-enabled' to yes,
# and optionally update the other syslog parameters to suit your needs.
# syslog-enabled no

# Specify the syslog identity.
# syslog-ident redis

# Specify the syslog facility. Must be USER or between LOCAL0-
LOCAL7.
# syslog-facility local0

# Set the number of databases. The default database is DB 0, you can
select
# a different one on a per-connection basis using SELECT <dbid> where
# dbid is a number between 0 and 'databases'-1
databases 16

################################## SNAPSHOTTING
################################
#
# Save the DB on disk:
#
#   save <seconds> <changes>
#
#   Will save the DB if both the given number of seconds and the given
#   number of write operations against the DB occurred.
#
#   In the example below the behaviour will be to save:
#   after 900 sec (15 min) if at least 1 key changed
#   after 300 sec (5 min) if at least 10 keys changed
#   after 60 sec if at least 10000 keys changed
#
#   Note: you can disable saving completely by commenting out all "save"
lines.
#
#   It is also possible to remove all the previously configured save
#   points by adding a save directive with a single empty string argument
#   like in the following example:
#
#   save ""

#save 900 1
```

```
#save 300 10
#save 60 10000
save ""

# By default Redis will stop accepting writes if RDB snapshots are
enabled
# (at least one save point) and the latest background save failed.
# This will make the user aware (in a hard way) that data is not persisting
# on disk properly, otherwise chances are that no one will notice and some
# disaster will happen.
#
# If the background saving process will start working again Redis will
# automatically allow writes again.
#
# However if you have setup your proper monitoring of the Redis server
# and persistence, you may want to disable this feature so that Redis will
# continue to work as usual even if there are problems with disk,
# permissions, and so forth.
stop-writes-on-bgsave-error yes

# Compress string objects using LZF when dump .rdb databases?
# For default that's set to 'yes' as it's almost always a win.
# If you want to save some CPU in the saving child set it to 'no' but
# the dataset will likely be bigger if you have compressible values or keys.
rdbcompression yes

# Since version 5 of RDB a CRC64 checksum is placed at the end of the
file.
# This makes the format more resistant to corruption but there is a
performance
# hit to pay (around 10%) when saving and loading RDB files, so you can
disable it
# for maximum performances.
#
# RDB files created with checksum disabled have a checksum of zero that
will
# tell the loading code to skip the check.
rdbchecksum yes

# The filename where to dump the DB
dbfilename dump.rdb

# The working directory.
#
# The DB will be written inside this directory, with the filename specified
# above using the 'dbfilename' configuration directive.
#
# The Append Only File will also be created inside this directory.
```

```
#
# Note that you must specify a directory here, not a file name.
dir /var/lib/redis

################################## REPLICATION
#################################

# Master-Slave replication. Use slaveof to make a Redis instance a copy
of
# another Redis server. A few things to understand ASAP about Redis
replication.
#
# 1) Redis replication is asynchronous, but you can configure a master to
#    stop accepting writes if it appears to be not connected with at least
#    a given number of slaves.
# 2) Redis slaves are able to perform a partial resynchronization with the
#    master if the replication link is lost for a relatively small amount of
#    time. You may want to configure the replication backlog size (see the
next
#    sections of this file) with a sensible value depending on your needs.
# 3) Replication is automatic and does not need user intervention. After a
#    network partition slaves automatically try to reconnect to masters
#    and resynchronize with them.
#
# slaveof <masterip> <masterport>

# If the master is password protected (using the "requirepass"
configuration
# directive below) it is possible to tell the slave to authenticate before
# starting the replication synchronization process, otherwise the master
will
# refuse the slave request.
#
# masterauth <master-password>

# When a slave loses its connection with the master, or when the
replication
# is still in progress, the slave can act in two different ways:
#
# 1) if slave-serve-stale-data is set to 'yes' (the default) the slave will
#    still reply to client requests, possibly with out of date data, or the
#    data set may just be empty if this is the first synchronization.
#
# 2) if slave-serve-stale-data is set to 'no' the slave will reply with
#    an error "SYNC with master in progress" to all the kind of commands
#    but to INFO and SLAVEOF.
#
slave-serve-stale-data yes
```

```
# You can configure a slave instance to accept writes or not. Writing against
# a slave instance may be useful to store some ephemeral data (because data
# written on a slave will be easily deleted after resync with the master) but
# may also cause problems if clients are writing to it because of a
# misconfiguration.
#
# Since Redis 2.6 by default slaves are read-only.
#
# Note: read only slaves are not designed to be exposed to untrusted clients
# on the internet. It's just a protection layer against misuse of the instance.
# Still a read only slave exports by default all the administrative commands
# such as CONFIG, DEBUG, and so forth. To a limited extent you can improve
# security of read only slaves using 'rename-command' to shadow all the
# administrative / dangerous commands.
slave-read-only yes

# Replication SYNC strategy: disk or socket.
#
# -------------------------------------------------------
# WARNING: DISKLESS REPLICATION IS EXPERIMENTAL CURRENTLY
# -------------------------------------------------------
#
# New slaves and reconnecting slaves that are not able to continue the replication
# process just receiving differences, need to do what is called a "full
# synchronization". An RDB file is transmitted from the master to the slaves.
# The transmission can happen in two different ways:
#
# 1) Disk-backed: The Redis master creates a new process that writes the RDB
#                 file on disk. Later the file is transferred by the parent
#                 process to the slaves incrementally.
# 2) Diskless: The Redis master creates a new process that directly writes the
#              RDB file to slave sockets, without touching the disk at all.
#
# With disk-backed replication, while the RDB file is generated, more slaves
# can be queued and served with the RDB file as soon as the current child producing
```

```
# the RDB file finishes its work. With diskless replication instead once
# the transfer starts, new slaves arriving will be queued and a new transfer
# will start when the current one terminates.
#
# When diskless replication is used, the master waits a configurable amount of
# time (in seconds) before starting the transfer in the hope that multiple slaves
# will arrive and the transfer can be parallelized.
#
# With slow disks and fast (large bandwidth) networks, diskless replication
# works better.
repl-diskless-sync no

# When diskless replication is enabled, it is possible to configure the delay
# the server waits in order to spawn the child that transfers the RDB via socket
# to the slaves.
#
# This is important since once the transfer starts, it is not possible to serve
# new slaves arriving, that will be queued for the next RDB transfer, so the server
# waits a delay in order to let more slaves arrive.
#
# The delay is specified in seconds, and by default is 5 seconds. To disable
# it entirely just set it to 0 seconds and the transfer will start ASAP.
repl-diskless-sync-delay 5

# Slaves send PINGs to server in a predefined interval. It's possible to change
# this interval with the repl_ping_slave_period option. The default value is 10
# seconds.
#
# repl-ping-slave-period 10

# The following option sets the replication timeout for:
#
# 1) Bulk transfer I/O during SYNC, from the point of view of slave.
# 2) Master timeout from the point of view of slaves (data, pings).
# 3) Slave timeout from the point of view of masters (REPLCONF ACK pings).
#
# It is important to make sure that this value is greater than the value
# specified for repl-ping-slave-period otherwise a timeout will be detected
# every time there is low traffic between the master and the slave.
```

```
#
# repl-timeout 60

# Disable TCP_NODELAY on the slave socket after SYNC?
#
# If you select "yes" Redis will use a smaller number of TCP packets and
# less bandwidth to send data to slaves. But this can add a delay for
# the data to appear on the slave side, up to 40 milliseconds with
# Linux kernels using a default configuration.
#
# If you select "no" the delay for data to appear on the slave side will
# be reduced but more bandwidth will be used for replication.
#
# By default we optimize for low latency, but in very high traffic
conditions
# or when the master and slaves are many hops away, turning this to "yes"
may
# be a good idea.
repl-disable-tcp-nodelay no

# Set the replication backlog size. The backlog is a buffer that
accumulates
# slave data when slaves are disconnected for some time, so that when a
slave
# wants to reconnect again, often a full resync is not needed, but a partial
# resync is enough, just passing the portion of data the slave missed while
# disconnected.
#
# The bigger the replication backlog, the longer the time the slave can be
# disconnected and later be able to perform a partial resynchronization.
#
# The backlog is only allocated once there is at least a slave connected.
#
# repl-backlog-size 1mb

# After a master has no longer connected slaves for some time, the
backlog
# will be freed. The following option configures the amount of seconds
that
# need to elapse, starting from the time the last slave disconnected, for
# the backlog buffer to be freed.
#
# A value of 0 means to never release the backlog.
#
# repl-backlog-ttl 3600

# The slave priority is an integer number published by Redis in the INFO
output.
```

```
# It is used by Redis Sentinel in order to select a slave to promote into a
# master if the master is no longer working correctly.
#
# A slave with a low priority number is considered better for promotion, so
# for instance if there are three slaves with priority 10, 100, 25 Sentinel will
# pick the one with priority 10, that is the lowest.
#
# However a special priority of 0 marks the slave as not able to perform the
# role of master, so a slave with priority of 0 will never be selected by
# Redis Sentinel for promotion.
#
# By default the priority is 100.
slave-priority 100

# It is possible for a master to stop accepting writes if there are less than
# N slaves connected, having a lag less or equal than M seconds.
#
# The N slaves need to be in "online" state.
#
# The lag in seconds, that must be <= the specified value, is calculated from
# the last ping received from the slave, that is usually sent every second.
#
# This option does not GUARANTEE that N replicas will accept the write, but
# will limit the window of exposure for lost writes in case not enough slaves
# are available, to the specified number of seconds.
#
# For example to require at least 3 slaves with a lag <= 10 seconds use:
#
# min-slaves-to-write 3
# min-slaves-max-lag 10
#
# Setting one or the other to 0 disables the feature.
#
# By default min-slaves-to-write is set to 0 (feature disabled) and
# min-slaves-max-lag is set to 10.

# A Redis master is able to list the address and port of the attached
# slaves in different ways. For example the "INFO replication" section
# offers this information, which is used, among other tools, by
# Redis Sentinel in order to discover slave instances.
# Another place where this info is available is in the output of the
# "ROLE" command of a masteer.
```

```
#
# The listed IP and address normally reported by a slave is obtained
# in the following way:
#
#   IP: The address is auto detected by checking the peer address
#   of the socket used by the slave to connect with the master.
#
#   Port: The port is communicated by the slave during the replication
#   handshake, and is normally the port that the slave is using to
#   list for connections.
#
# However when port forwarding or Network Address Translation (NAT)
is
# used, the slave may be actually reachable via different IP and port
# pairs. The following two options can be used by a slave in order to
# report to its master a specific set of IP and port, so that both INFO
# and ROLE will report those values.
#
# There is no need to use both the options if you need to override just
# the port or the IP address.
#
# slave-announce-ip 5.5.5.5
# slave-announce-port 1234

################################## SECURITY
###################################

# Require clients to issue AUTH <PASSWORD> before processing any
other
# commands.  This might be useful in environments in which you do not
trust
# others with access to the host running redis-server.
#
# This should stay commented out for backward compatibility and
because most
# people do not need auth (e.g. they run their own servers).
#
# Warning: since Redis is pretty fast an outside user can try up to
# 150k passwords per second against a good box. This means that you
should
# use a very strong password otherwise it will be very easy to break.
#
# requirepass foobared

# Command renaming.
#
# It is possible to change the name of dangerous commands in a shared
```

```
# environment. For instance the CONFIG command may be renamed into
something
# hard to guess so that it will still be available for internal-use tools
# but not available for general clients.
#
# Example:
#
# rename-command CONFIG
b840fc02d524045429941cc15f59e41cb7be6c52
#
# It is also possible to completely kill a command by renaming it into
# an empty string:
#
# rename-command CONFIG ""
#
# Please note that changing the name of commands that are logged into
the
# AOF file or transmitted to slaves may cause problems.

################################### LIMITS
####################################

# Set the max number of connected clients at the same time. By default
# this limit is set to 10000 clients, however if the Redis server is not
# able to configure the process file limit to allow for the specified limit
# the max number of allowed clients is set to the current file limit
# minus 32 (as Redis reserves a few file descriptors for internal uses).
#
# Once the limit is reached Redis will close all the new connections
sending
# an error 'max number of clients reached'.
#
# maxclients 10000

# Don't use more memory than the specified amount of bytes.
# When the memory limit is reached Redis will try to remove keys
# according to the eviction policy selected (see maxmemory-policy).
#
# If Redis can't remove keys according to the policy, or if the policy is
# set to 'noeviction', Redis will start to reply with errors to commands
# that would use more memory, like SET, LPUSH, and so on, and will
continue
# to reply to read-only commands like GET.
#
# This option is usually useful when using Redis as an LRU cache, or to
set
# a hard memory limit for an instance (using the 'noeviction' policy).
#
```

```
# WARNING: If you have slaves attached to an instance with
maxmemory on,
# the size of the output buffers needed to feed the slaves are subtracted
# from the used memory count, so that network problems / resyncs will
# not trigger a loop where keys are evicted, and in turn the output
# buffer of slaves is full with DELs of keys evicted triggering the deletion
# of more keys, and so forth until the database is completely emptied.
#
# In short... if you have slaves attached it is suggested that you set a lower
# limit for maxmemory so that there is some free RAM on the system for
slave
# output buffers (but this is not needed if the policy is 'noeviction').
#
# maxmemory <bytes>

# MAXMEMORY POLICY: how Redis will select what to remove when
maxmemory
# is reached. You can select among five behaviors:
#
# volatile-lru -> remove the key with an expire set using an LRU
algorithm
# allkeys-lru -> remove any key according to the LRU algorithm
# volatile-random -> remove a random key with an expire set
# allkeys-random -> remove a random key, any key
# volatile-ttl -> remove the key with the nearest expire time (minor TTL)
# noeviction -> don't expire at all, just return an error on write operations
#
# Note: with any of the above policies, Redis will return an error on write
#       operations, when there are no suitable keys for eviction.
#
#       At the date of writing these commands are: set setnx setex append
#       incr decr rpush lpush rpushx lpushx linsert lset rpoplpush sadd
#       sinter sinterstore sunion sunionstore sdiff sdiffstore zadd zincrby
#       zunionstore zinterstore hset hsetnx hmset hincrby incrby decrby
#       getset mset msetnx exec sort
#
# The default is:
#
# maxmemory-policy noeviction

# LRU and minimal TTL algorithms are not precise algorithms but
approximated
# algorithms (in order to save memory), so you can tune it for speed or
# accuracy. For default Redis will check five keys and pick the one that
was
# used less recently, you can change the sample size using the following
# configuration directive.
#
```

```
# The default of 5 produces good enough results. 10 Approximates very
closely
# true LRU but costs a bit more CPU. 3 is very fast but not very accurate.
#
# maxmemory-samples 5

################################ APPEND ONLY MODE
################################

# By default Redis asynchronously dumps the dataset on disk. This mode
is
# good enough in many applications, but an issue with the Redis process
or
# a power outage may result into a few minutes of writes lost (depending
on
# the configured save points).
#
# The Append Only File is an alternative persistence mode that provides
# much better durability. For instance using the default data fsync policy
# (see later in the config file) Redis can lose just one second of writes in a
# dramatic event like a server power outage, or a single write if something
# wrong with the Redis process itself happens, but the operating system is
# still running correctly.
#
# AOF and RDB persistence can be enabled at the same time without
problems.
# If the AOF is enabled on startup Redis will load the AOF, that is the file
# with the better durability guarantees.
#
# Please check http://redis.io/topics/persistence for more information.

appendonly no

# The name of the append only file (default: "appendonly.aof")

appendfilename "appendonly.aof"

# The fsync() call tells the Operating System to actually write data on disk
# instead of waiting for more data in the output buffer. Some OS will
really flush
# data on disk, some other OS will just try to do it ASAP.
#
# Redis supports three different modes:
#
# no: don't fsync, just let the OS flush the data when it wants. Faster.
# always: fsync after every write to the append only log. Slow, Safest.
# everysec: fsync only one time every second. Compromise.
#
```

```
# The default is "everysec", as that's usually the right compromise
between
# speed and data safety. It's up to you to understand if you can relax this to
# "no" that will let the operating system flush the output buffer when
# it wants, for better performances (but if you can live with the idea of
# some data loss consider the default persistence mode that's
snapshotting),
# or on the contrary, use "always" that's very slow but a bit safer than
# everysec.
#
# More details please check the following article:
# http://antirez.com/post/redis-persistence-demystified.html
#
# If unsure, use "everysec".

# appendfsync always
appendfsync everysec
# appendfsync no

# When the AOF fsync policy is set to always or everysec, and a
background
# saving process (a background save or AOF log background rewriting) is
# performing a lot of I/O against the disk, in some Linux configurations
# Redis may block too long on the fsync() call. Note that there is no fix
for
# this currently, as even performing fsync in a different thread will block
# our synchronous write(2) call.
#
# In order to mitigate this problem it's possible to use the following option
# that will prevent fsync() from being called in the main process while a
# BGSAVE or BGREWRITEAOF is in progress.
#
# This means that while another child is saving, the durability of Redis is
# the same as "appendfsync none". In practical terms, this means that it is
# possible to lose up to 30 seconds of log in the worst scenario (with the
# default Linux settings).
#
# If you have latency problems turn this to "yes". Otherwise leave it as
# "no" that is the safest pick from the point of view of durability.

no-appendfsync-on-rewrite no

# Automatic rewrite of the append only file.
# Redis is able to automatically rewrite the log file implicitly calling
# BGREWRITEAOF when the AOF log size grows by the specified
percentage.
#
# This is how it works: Redis remembers the size of the AOF file after the
```

```
# latest rewrite (if no rewrite has happened since the restart, the size of
# the AOF at startup is used).
#
# This base size is compared to the current size. If the current size is
# bigger than the specified percentage, the rewrite is triggered. Also
# you need to specify a minimal size for the AOF file to be rewritten, this
# is useful to avoid rewriting the AOF file even if the percentage increase
# is reached but it is still pretty small.
#
# Specify a percentage of zero in order to disable the automatic AOF
# rewrite feature.

auto-aof-rewrite-percentage 100
auto-aof-rewrite-min-size 64mb

# An AOF file may be found to be truncated at the end during the Redis
# startup process, when the AOF data gets loaded back into memory.
# This may happen when the system where Redis is running
# crashes, especially when an ext4 filesystem is mounted without the
# data=ordered option (however this can't happen when Redis itself
# crashes or aborts but the operating system still works correctly).
#
# Redis can either exit with an error when this happens, or load as much
# data as possible (the default now) and start if the AOF file is found
# to be truncated at the end. The following option controls this behavior.
#
# If aof-load-truncated is set to yes, a truncated AOF file is loaded and
# the Redis server starts emitting a log to inform the user of the event.
# Otherwise if the option is set to no, the server aborts with an error
# and refuses to start. When the option is set to no, the user requires
# to fix the AOF file using the "redis-check-aof" utility before to restart
# the server.
#
# Note that if the AOF file will be found to be corrupted in the middle
# the server will still exit with an error. This option only applies when
# Redis will try to read more data from the AOF file but not enough bytes
# will be found.
aof-load-truncated yes

################################ LUA SCRIPTING
################################

# Max execution time of a Lua script in milliseconds.
#
# If the maximum execution time is reached Redis will log that a script is
# still in execution after the maximum allowed time and will start to
# reply to queries with an error.
#
```

```
# When a long running script exceeds the maximum execution time only
the
# SCRIPT KILL and SHUTDOWN NOSAVE commands are available.
The first can be
# used to stop a script that did not yet called write commands. The second
# is the only way to shut down the server in the case a write command
was
# already issued by the script but the user doesn't want to wait for the
natural
# termination of the script.
#
# Set it to 0 or a negative value for unlimited execution without warnings.
lua-time-limit 5000

################################ REDIS CLUSTER
################################
#
#
++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
++++++++++++++++++++++++
# WARNING EXPERIMENTAL: Redis Cluster is considered to be stable
code, however
# in order to mark it as "mature" we need to wait for a non trivial
percentage
# of users to deploy it in production.
#
++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
++++++++++++++++++++++++
#
# Normal Redis instances can't be part of a Redis Cluster; only nodes that
are
# started as cluster nodes can. In order to start a Redis instance as a
# cluster node enable the cluster support uncommenting the following:
#
# cluster-enabled yes

# Every cluster node has a cluster configuration file. This file is not
# intended to be edited by hand. It is created and updated by Redis nodes.
# Every Redis Cluster node requires a different cluster configuration file.
# Make sure that instances running in the same system do not have
# overlapping cluster configuration file names.
#
# cluster-config-file nodes-6379.conf

# Cluster node timeout is the amount of milliseconds a node must be
unreachable
# for it to be considered in failure state.
# Most other internal time limits are multiple of the node timeout.
```

```
#
# cluster-node-timeout 15000

# A slave of a failing master will avoid to start a failover if its data
# looks too old.
#
# There is no simple way for a slave to actually have a exact measure of
# its "data age", so the following two checks are performed:
#
# 1) If there are multiple slaves able to failover, they exchange messages
#    in order to try to give an advantage to the slave with the best
#    replication offset (more data from the master processed).
#    Slaves will try to get their rank by offset, and apply to the start
#    of the failover a delay proportional to their rank.
#
# 2) Every single slave computes the time of the last interaction with
#    its master. This can be the last ping or command received (if the
master
#    is still in the "connected" state), or the time that elapsed since the
#    disconnection with the master (if the replication link is currently
down).
#    If the last interaction is too old, the slave will not try to failover
#    at all.
#
# The point "2" can be tuned by user. Specifically a slave will not perform
# the failover if, since the last interaction with the master, the time
# elapsed is greater than:
#
#   (node-timeout * slave-validity-factor) + repl-ping-slave-period
#
# So for example if node-timeout is 30 seconds, and the slave-validity-
factor
# is 10, and assuming a default repl-ping-slave-period of 10 seconds, the
# slave will not try to failover if it was not able to talk with the master
# for longer than 310 seconds.
#
# A large slave-validity-factor may allow slaves with too old data to
failover
# a master, while a too small value may prevent the cluster from being
able to
# elect a slave at all.
#
# For maximum availability, it is possible to set the slave-validity-factor
# to a value of 0, which means, that slaves will always try to failover the
# master regardless of the last time they interacted with the master.
# (However they'll always try to apply a delay proportional to their
# offset rank).
#
```

```
# Zero is the only value able to guarantee that when all the partitions heal
# the cluster will always be able to continue.
#
# cluster-slave-validity-factor 10

# Cluster slaves are able to migrate to orphaned masters, that are masters
# that are left without working slaves. This improves the cluster ability
# to resist to failures as otherwise an orphaned master can't be failed over
# in case of failure if it has no working slaves.
#
# Slaves migrate to orphaned masters only if there are still at least a
# given number of other working slaves for their old master. This number
# is the "migration barrier". A migration barrier of 1 means that a slave
# will migrate only if there is at least 1 other working slave for its master
# and so forth. It usually reflects the number of slaves you want for every
# master in your cluster.
#
# Default is 1 (slaves migrate only if their masters remain with at least
# one slave). To disable migration just set it to a very large value.
# A value of 0 can be set but is useful only for debugging and dangerous
# in production.
#
# cluster-migration-barrier 1

# By default Redis Cluster nodes stop accepting queries if they detect
there
# is at least an hash slot uncovered (no available node is serving it).
# This way if the cluster is partially down (for example a range of hash
slots
# are no longer covered) all the cluster becomes, eventually, unavailable.
# It automatically returns available as soon as all the slots are covered
again.
#
# However sometimes you want the subset of the cluster which is
working,
# to continue to accept queries for the part of the key space that is still
# covered. In order to do so, just set the cluster-require-full-coverage
# option to no.
#
# cluster-require-full-coverage yes

# In order to setup your cluster make sure to read the documentation
# available at http://redis.io web site.

################################## SLOW LOG
###################################

# The Redis Slow Log is a system to log queries that exceeded a specified
```

```
# execution time. The execution time does not include the I/O operations
# like talking with the client, sending the reply and so forth,
# but just the time needed to actually execute the command (this is the
only
# stage of command execution where the thread is blocked and can not
serve
# other requests in the meantime).
#
# You can configure the slow log with two parameters: one tells Redis
# what is the execution time, in microseconds, to exceed in order for the
# command to get logged, and the other parameter is the length of the
# slow log. When a new command is logged the oldest one is removed
from the
# queue of logged commands.

# The following time is expressed in microseconds, so 1000000 is
equivalent
# to one second. Note that a negative number disables the slow log, while
# a value of zero forces the logging of every command.
slowlog-log-slower-than 10000

# There is no limit to this length. Just be aware that it will consume
memory.
# You can reclaim memory used by the slow log with SLOWLOG
RESET.
slowlog-max-len 128

################################ LATENCY MONITOR
##############################

# The Redis latency monitoring subsystem samples different operations
# at runtime in order to collect data related to possible sources of
# latency of a Redis instance.
#
# Via the LATENCY command this information is available to the user
that can
# print graphs and obtain reports.
#
# The system only logs operations that were performed in a time equal or
# greater than the amount of milliseconds specified via the
# latency-monitor-threshold configuration directive. When its value is set
# to zero, the latency monitor is turned off.
#
# By default latency monitoring is disabled since it is mostly not needed
# if you don't have latency issues, and collecting data has a performance
# impact, that while very small, can be measured under big load. Latency
# monitoring can easily be enabled at runtime using the command
# "CONFIG SET latency-monitor-threshold <milliseconds>" if needed.
```

```
latency-monitor-threshold 0

############################## EVENT NOTIFICATION ##############################

# Redis can notify Pub/Sub clients about events happening in the key space.
# This feature is documented at http://redis.io/topics/notifications
#
# For instance if keyspace events notification is enabled, and a client
# performs a DEL operation on key "foo" stored in the Database 0, two
# messages will be published via Pub/Sub:
#
# PUBLISH __keyspace@0__:foo del
# PUBLISH __keyevent@0__:del foo
#
# It is possible to select the events that Redis will notify among a set
# of classes. Every class is identified by a single character:
#
#  K     Keyspace events, published with __keyspace@<db>__ prefix.
#  E     Keyevent events, published with __keyevent@<db>__ prefix.
#  g     Generic commands (non-type specific) like DEL, EXPIRE, RENAME, ...
#  $     String commands
#  l     List commands
#  s     Set commands
#  h     Hash commands
#  z     Sorted set commands
#  x     Expired events (events generated every time a key expires)
#  e     Evicted events (events generated when a key is evicted for maxmemory)
#  A     Alias for g$lshzxe, so that the "AKE" string means all the events.
#
#  The "notify-keyspace-events" takes as argument a string that is composed
#  of zero or multiple characters. The empty string means that notifications
#  are disabled.
#
#  Example: to enable list and generic events, from the point of view of the
#           event name, use:
#
#  notify-keyspace-events Elg
#
#  Example 2: to get the stream of the expired keys subscribing to channel
#             name __keyevent@0__:expired use:
#
```

```
#  notify-keyspace-events Ex
#
#  By default all notifications are disabled because most users don't need
#  this feature and the feature has some overhead. Note that if you don't
#  specify at least one of K or E, no events will be delivered.
notify-keyspace-events ""

################################ ADVANCED CONFIG
################################

# Hashes are encoded using a memory efficient data structure when they
have a
# small number of entries, and the biggest entry does not exceed a given
# threshold. These thresholds can be configured using the following
directives.
hash-max-ziplist-entries 512
hash-max-ziplist-value 64

# Lists are also encoded in a special way to save a lot of space.
# The number of entries allowed per internal list node can be specified
# as a fixed maximum size or a maximum number of elements.
# For a fixed maximum size, use -5 through -1, meaning:
# -5: max size: 64 Kb  <-- not recommended for normal workloads
# -4: max size: 32 Kb  <-- not recommended
# -3: max size: 16 Kb  <-- probably not recommended
# -2: max size: 8 Kb   <-- good
# -1: max size: 4 Kb   <-- good
# Positive numbers mean store up to _exactly_ that number of elements
# per list node.
# The highest performing option is usually -2 (8 Kb size) or -1 (4 Kb
size),
# but if your use case is unique, adjust the settings as necessary.
list-max-ziplist-size -2

# Lists may also be compressed.
# Compress depth is the number of quicklist ziplist nodes from *each*
side of
# the list to *exclude* from compression.  The head and tail of the list
# are always uncompressed for fast push/pop operations.  Settings are:
# 0: disable all list compression
# 1: depth 1 means "don't start compressing until after 1 node into the list,
#    going from either the head or tail"
#    So: [head]->node->node->...->node->[tail]
#    [head], [tail] will always be uncompressed; inner nodes will compress.
# 2: [head]->[next]->node->node->...->node->[prev]->[tail]
#    2 here means: don't compress head or head->next or tail->prev or tail,
#    but compress all nodes between them.
```

```
# 3: [head]->[next]->[next]->node->node->...->node->[prev]->[prev]->[tail]
# etc.
list-compress-depth 0

# Sets have a special encoding in just one case: when a set is composed
# of just strings that happen to be integers in radix 10 in the range
# of 64 bit signed integers.
# The following configuration setting sets the limit in the size of the
# set in order to use this special memory saving encoding.
set-max-intset-entries 512

# Similarly to hashes and lists, sorted sets are also specially encoded in
# order to save a lot of space. This encoding is only used when the length and
# elements of a sorted set are below the following limits:
zset-max-ziplist-entries 128
zset-max-ziplist-value 64

# HyperLogLog sparse representation bytes limit. The limit includes the
# 16 bytes header. When an HyperLogLog using the sparse representation crosses
# this limit, it is converted into the dense representation.
#
# A value greater than 16000 is totally useless, since at that point the
# dense representation is more memory efficient.
#
# The suggested value is ~ 3000 in order to have the benefits of
# the space efficient encoding without slowing down too much PFADD,
# which is O(N) with the sparse encoding. The value can be raised to
# ~ 10000 when CPU is not a concern, but space is, and the data set is
# composed of many HyperLogLogs with cardinality in the 0 - 15000 range.
hll-sparse-max-bytes 3000

# Active rehashing uses 1 millisecond every 100 milliseconds of CPU time in
# order to help rehashing the main Redis hash table (the one mapping top-level
# keys to values). The hash table implementation Redis uses (see dict.c)
# performs a lazy rehashing: the more operation you run into a hash table
# that is rehashing, the more rehashing "steps" are performed, so if the
# server is idle the rehashing is never complete and some more memory is used
# by the hash table.
#
# The default is to use this millisecond 10 times every second in order to
# actively rehash the main dictionaries, freeing memory when possible.
```

```
#
# If unsure:
# use "activerehashing no" if you have hard latency requirements and it is
# not a good thing in your environment that Redis can reply from time to
time
# to queries with 2 milliseconds delay.
#
# use "activerehashing yes" if you don't have such hard requirements but
# want to free memory asap when possible.
activerehashing yes

# The client output buffer limits can be used to force disconnection of
clients
# that are not reading data from the server fast enough for some reason (a
# common reason is that a Pub/Sub client can't consume messages as fast
as the
# publisher can produce them).
#
# The limit can be set differently for the three different classes of clients:
#
# normal -> normal clients including MONITOR clients
# slave  -> slave clients
# pubsub -> clients subscribed to at least one pubsub channel or pattern
#
# The syntax of every client-output-buffer-limit directive is the following:
#
# client-output-buffer-limit <class> <hard limit> <soft limit> <soft
seconds>
#
# A client is immediately disconnected once the hard limit is reached, or if
# the soft limit is reached and remains reached for the specified number of
# seconds (continuously).
# So for instance if the hard limit is 32 megabytes and the soft limit is
# 16 megabytes / 10 seconds, the client will get disconnected immediately
# if the size of the output buffers reach 32 megabytes, but will also get
# disconnected if the client reaches 16 megabytes and continuously
overcomes
# the limit for 10 seconds.
#
# By default normal clients are not limited because they don't receive data
# without asking (in a push way), but just after a request, so only
# asynchronous clients may create a scenario where data is requested
faster
# than it can read.
#
# Instead there is a default limit for pubsub and slave clients, since
# subscribers and slaves receive data in a push fashion.
#
```

	# Both the hard or the soft limit can be disabled by setting them to zero. client-output-buffer-limit normal 0 0 0 client-output-buffer-limit slave 256mb 64mb 60 client-output-buffer-limit pubsub 32mb 8mb 60 # Redis calls an internal function to perform many background tasks, like # closing connections of clients in timeout, purging expired keys that are # never requested, and so forth. # # Not all tasks are performed with the same frequency, but Redis checks for # tasks to perform according to the specified "hz" value. # # By default "hz" is set to 10. Raising the value will use more CPU when # Redis is idle, but at the same time will make Redis more responsive when # there are many keys expiring at the same time, and timeouts may be # handled with more precision. # # The range is between 1 and 500, however a value over 100 is usually not # a good idea. Most users should use the default of 10 and raise this up to # 100 only in environments where very low latency is required. hz 10 # When a child rewrites the AOF file, if the following option is enabled # the file will be fsync-ed every 32 MB of data generated. This is useful # in order to commit the file to the disk more incrementally and avoid # big latency spikes. aof-rewrite-incremental-fsync yes

I. Nazwa serwera: cassandra1

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	cassandra1
IP / Domena	10.90.54.154
OS	CentOS 7
RAM:	8 GB
vCPU	4 vCPU
Przestrzeń	100 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Bazodanowy

2. Konfiguracja .hosts

```
127.0.0.1 localhost
10.90.54.151 ku1
10.90.54.152 ku2
10.90.54.153 ku3
10.90.54.154 ku4
```

3. Interfejsy sieciowe

eth0 10.90.54.154

4. Zainstalowane usługi

a. CassandraDB

Parametr	Wartość
Lokalizacja instalacji	/etc/cassandra
Zawartość pliku konfiguracyjnego #1	/etc/cassandra/conf # Cassandra storage config YAML. # NOTE: # See http://wiki.apache.org/cassandra/StorageConfiguration for # full explanations of configuration directives # /NOTE # The name of the cluster. This is mainly used to prevent machines in # one logical cluster from joining another. cluster_name: 'EpoCluster' # This defines the number of tokens randomly assigned to this node on the ring # The more tokens, relative to other nodes, the larger the proportion of data # that this node will store. You probably want all nodes to have the same number # of tokens assuming they have equal hardware capability. # # If you leave this unspecified, Cassandra will use the default of 1 token for legacy compatibility, # and will use the initial_token as described below. #

```
# Specifying initial_token will override this setting on the node's initial start,
# on subsequent starts, this setting will apply even if initial token is set.
#
# If you already have a cluster with 1 token per node, and wish to migrate to.
# multiple tokens per node, see http://wiki.apache.org/cassandra/Operations
num_tokens: 256

# initial_token allows you to specify tokens manually.  While you can use # it with
# vnodes (num_tokens > 1, above) -- in which case you should provide a.
# comma-separated list -- it's primarily used when adding nodes # to legacy clusters.
# that do not have vnodes enabled.
# initial_token:

# See http://wiki.apache.org/cassandra/HintedHandoff
# May either be "true" or "false" to enable globally, or contain a list
# of data centers to enable per-datacenter.
# hinted_handoff_enabled: DC1,DC2
hinted_handoff_enabled: true
# this defines the maximum amount of time a dead host will have hints
# generated.  After it has been dead this long, new hints for it will not be
# created until it has been seen alive and gone down again.
max_hint_window_in_ms: 10800000 # 3 hours
# Maximum throttle in KBs per second, per delivery thread.  This will be
# reduced proportionally to the number of nodes in the cluster.  (If there
# are two nodes in the cluster, each delivery thread will use the maximum
# rate; if there are three, each will throttle to half of the maximum,
# since we expect two nodes to be delivering hints simultaneously.)
hinted_handoff_throttle_in_kb: 1024
# Number of threads with which to deliver hints;
# Consider increasing this number when you have multi-dc deployments, since
# cross-dc handoff tends to be slower
max_hints_delivery_threads: 2

# Maximum throttle in KBs per second, total. This will be
# reduced proportionally to the number of nodes in the cluster.
batchlog_replay_throttle_in_kb: 1024

# Authentication backend, implementing IAuthenticator; used to identify users
# Out of the box, Cassandra provides org.apache.cassandra.auth.{AllowAllAuthenticator,
# PasswordAuthenticator}.
#
# - AllowAllAuthenticator performs no checks - set it to disable authentication.
# - PasswordAuthenticator relies on username/password pairs to authenticate
#   users. It keeps usernames and hashed passwords in system_auth.credentials table.
#   Please increase system_auth keyspace replication factor if you use this authenticator.
authenticator: PasswordAuthenticator

# Authorization backend, implementing IAuthorizer; used to limit access/provide permissions
# Out of the box, Cassandra provides org.apache.cassandra.auth.{AllowAllAuthorizer,
# CassandraAuthorizer}.
#
# - AllowAllAuthorizer allows any action to any user - set it to disable authorization.
# - CassandraAuthorizer stores permissions in system_auth.permissions table. Please
#   increase system_auth keyspace replication factor if you use this authorizer.
authorizer: CassandraAuthorizer

# Validity period for permissions cache (fetching permissions can be an
# expensive operation depending on the authorizer, CassandraAuthorizer is
# one example). Defaults to 2000, set to 0 to disable.
# Will be disabled automatically for AllowAllAuthorizer.
permissions_validity_in_ms: 2000

# Refresh interval for permissions cache (if enabled).
# After this interval, cache entries become eligible for refresh. Upon next
# access, an async reload is scheduled and the old value returned until it
# completes. If permissions_validity_in_ms is non-zero, then this must be
# also.
# Defaults to the same value as permissions_validity_in_ms.
# permissions_update_interval_in_ms: 1000

# The partitioner is responsible for distributing groups of rows (by
# partition key) across nodes in the cluster.  You should leave this
# alone for new clusters.  The partitioner can NOT be changed without
```

```
# reloading all data, so when upgrading, you should set this to the
# same partitioner you were already using.
#
# Besides Murmur3Partitioner, partitioners included for backwards
# compatibility include RandomPartitioner, ByteOrderedPartitioner, and
# OrderPreservingPartitioner.
#
partitioner: org.apache.cassandra.dht.Murmur3Partitioner

# Directories where Cassandra should store data on disk.  Cassandra
# will spread data evenly across them, subject to the granularity of
# the configured compaction strategy.
# If not set, the default directory is $CASSANDRA_HOME/data/data.
data_file_directories:
    - /opt/cassandra/data

# commit log.  when running on magnetic HDD, this should be a
# separate spindle than the data directories.
# If not set, the default directory is $CASSANDRA_HOME/data/commitlog.
commitlog_directory: /opt/cassandra/commitlog

# policy for data disk failures:
# die: shut down gossip and client transports and kill the JVM for any fs errors or
#      single-sstable errors, so the node can be replaced.
# stop_paranoid: shut down gossip and client transports even for single-sstable errors,
#                kill the JVM for errors during startup.
# stop: shut down gossip and client transports, leaving the node effectively dead, but
#       can still be inspected via JMX, kill the JVM for errors during startup.
# best_effort: stop using the failed disk and respond to requests based on
#              remaining available sstables.  This means you WILL see obsolete
#              data at CL.ONE!
# ignore: ignore fatal errors and let requests fail, as in pre-1.2 Cassandra
disk_failure_policy: stop

# policy for commit disk failures:
# die: shut down gossip and Thrift and kill the JVM, so the node can be replaced.
# stop: shut down gossip and Thrift, leaving the node effectively dead, but
#       can still be inspected via JMX.
# stop_commit: shutdown the commit log, letting writes collect but
#              continuing to service reads, as in pre-2.0.5 Cassandra
# ignore: ignore fatal errors and let the batches fail
commit_failure_policy: stop

# Maximum size of the key cache in memory.
#
# Each key cache hit saves 1 seek and each row cache hit saves 2 seeks at the
# minimum, sometimes more. The key cache is fairly tiny for the amount of
# time it saves, so it's worthwhile to use it at large numbers.
# The row cache saves even more time, but must contain the entire row,
# so it is extremely space-intensive. It's best to only use the
# row cache if you have hot rows or static rows.
#
# NOTE: if you reduce the size, you may not get you hottest keys loaded on startup.
#
# Default value is empty to make it "auto" (min(5% of Heap (in MB), 100MB)). Set to 0 to disable key cache.
key_cache_size_in_mb:

# Duration in seconds after which Cassandra should
# save the key cache. Caches are saved to saved_caches_directory as
# specified in this configuration file.
#
# Saved caches greatly improve cold-start speeds, and is relatively cheap in
# terms of I/O for the key cache. Row cache saving is much more expensive and
# has limited use.
#
# Default is 14400 or 4 hours.
key_cache_save_period: 14400

# Number of keys from the key cache to save
# Disabled by default, meaning all keys are going to be saved
# key_cache_keys_to_save: 100

# Maximum size of the row cache in memory.
```

```
# NOTE: if you reduce the size, you may not get you hottest keys loaded on startup.
#
# Default value is 0, to disable row caching.
row_cache_size_in_mb: 0

# Duration in seconds after which Cassandra should
# save the row cache. Caches are saved to saved_caches_directory as specified
# in this configuration file.
#
# Saved caches greatly improve cold-start speeds, and is relatively cheap in
# terms of I/O for the key cache. Row cache saving is much more expensive and
# has limited use.
#
# Default is 0 to disable saving the row cache.
row_cache_save_period: 0

# Number of keys from the row cache to save
# Disabled by default, meaning all keys are going to be saved
# row_cache_keys_to_save: 100

# Maximum size of the counter cache in memory.
#
# Counter cache helps to reduce counter locks' contention for hot counter cells.
# In case of RF = 1 a counter cache hit will cause Cassandra to skip the read before
# write entirely. With RF > 1 a counter cache hit will still help to reduce the duration
# of the lock hold, helping with hot counter cell updates, but will not allow skipping
# the read entirely. Only the local (clock, count) tuple of a counter cell is kept
# in memory, not the whole counter, so it's relatively cheap.
#
# NOTE: if you reduce the size, you may not get you hottest keys loaded on startup.
#
# Default value is empty to make it "auto" (min(2.5% of Heap (in MB), 50MB)). Set to 0 to disable counter cache.
# NOTE: if you perform counter deletes and rely on low gcgs, you should disable the counter cache.
counter_cache_size_in_mb:

# Duration in seconds after which Cassandra should
# save the counter cache (keys only). Caches are saved to saved_caches_directory as
# specified in this configuration file.
#
# Default is 7200 or 2 hours.
counter_cache_save_period: 7200

# Number of keys from the counter cache to save
# Disabled by default, meaning all keys are going to be saved
# counter_cache_keys_to_save: 100

# The off-heap memory allocator.  Affects storage engine metadata as
# well as caches.  Experiments show that JEMAlloc saves some memory
# than the native GCC allocator (i.e., JEMalloc is more
# fragmentation-resistant).
#.
# Supported values are: NativeAllocator, JEMallocAllocator
#
# If you intend to use JEMallocAllocator you have to install JEMalloc as library and
# modify cassandra-env.sh as directed in the file.
#
# Defaults to NativeAllocator
# memory_allocator: NativeAllocator

# saved caches
# If not set, the default directory is $CASSANDRA_HOME/data/saved_caches.
saved_caches_directory: /opt/cassandra/saved_caches

# commitlog_sync may be either "periodic" or "batch.".
#.
# When in batch mode, Cassandra won't ack writes until the commit log
# has been fsynced to disk.  It will wait
# commitlog_sync_batch_window_in_ms milliseconds between fsyncs.
# This window should be kept short because the writer threads will
# be unable to do extra work while waiting.  (You may need to increase
# concurrent_writes for the same reason.)
#
# commitlog_sync: batch
```

```
# commitlog_sync_batch_window_in_ms: 2
#
# the other option is "periodic" where writes may be acked immediately
# and the CommitLog is simply synced every commitlog_sync_period_in_ms
# milliseconds..
commitlog_sync: periodic
commitlog_sync_period_in_ms: 10000

# The size of the individual commitlog file segments.  A commitlog
# segment may be archived, deleted, or recycled once all the data
# in it (potentially from each columnfamily in the system) has been
# flushed to sstables...
#
# The default size is 32, which is almost always fine, but if you are
# archiving commitlog segments (see commitlog_archiving.properties),
# then you probably want a finer granularity of archiving; 8 or 16 MB
# is reasonable.
commitlog_segment_size_in_mb: 32

# Reuse commit log files when possible. The default is false, and this
# feature will be removed entirely in future versions of Cassandra.
#commitlog_segment_recycling: false

# any class that implements the SeedProvider interface and has a
# constructor that takes a Map<String, String> of parameters will do.
seed_provider:
   # Addresses of hosts that are deemed contact points..
   # Cassandra nodes use this list of hosts to find each other and learn
   # the topology of the ring.  You must change this if you are running
   # multiple nodes!
   - class_name: org.apache.cassandra.locator.SimpleSeedProvider
     parameters:
         # seeds is actually a comma-delimited list of addresses.
         # Ex: "<ip1>,<ip2>,<ip3>"
         - seeds: "10.90.54.151,10.90.54.152,10.90.54.153,10.90.54.154"

# For workloads with more data than can fit in memory, Cassandra's
# bottleneck will be reads that need to fetch data from
# disk. "concurrent_reads" should be set to (16 * number_of_drives) in
# order to allow the operations to enqueue low enough in the stack
# that the OS and drives can reorder them. Same applies to
# "concurrent_counter_writes", since counter writes read the current
# values before incrementing and writing them back.
#
# On the other hand, since writes are almost never IO bound, the ideal
# number of "concurrent_writes" is dependent on the number of cores in
# your system; (8 * number_of_cores) is a good rule of thumb.
concurrent_reads: 32
concurrent_writes: 32
concurrent_counter_writes: 32

# Total memory to use for sstable-reading buffers.  Defaults to
# the smaller of 1/4 of heap or 512MB.
# file_cache_size_in_mb: 512

# Total permitted memory to use for memtables. Cassandra will stop.
# accepting writes when the limit is exceeded until a flush completes,
# and will trigger a flush based on memtable_cleanup_threshold
# If omitted, Cassandra will set both to 1/4 the size of the heap.
# memtable_heap_space_in_mb: 2048
# memtable_offheap_space_in_mb: 2048

# Ratio of occupied non-flushing memtable size to total permitted size
# that will trigger a flush of the largest memtable.  Lager mct will
# mean larger flushes and hence less compaction, but also less concurrent
# flush activity which can make it difficult to keep your disks fed
# under heavy write load.
#
# memtable_cleanup_threshold defaults to 1 / (memtable_flush_writers + 1)
# memtable_cleanup_threshold: 0.11

# Specify the way Cassandra allocates and manages memtable memory.
# Options are:
```

```
#   heap_buffers:    on heap nio buffers
#   offheap_buffers: off heap (direct) nio buffers
#   offheap_objects: native memory, eliminating nio buffer heap overhead
memtable_allocation_type: heap_buffers

# Total space to use for commitlogs.  Since commitlog segments are
# mmapped, and hence use up address space, the default size is 32
# on 32-bit JVMs, and 8192 on 64-bit JVMs.
#
# If space gets above this value (it will round up to the next nearest
# segment multiple), Cassandra will flush every dirty CF in the oldest
# segment and remove it.  So a small total commitlog space will tend
# to cause more flush activity on less-active columnfamilies.
# commitlog_total_space_in_mb: 8192

# This sets the amount of memtable flush writer threads.  These will
# be blocked by disk io, and each one will hold a memtable in memory
# while blocked..
#
# memtable_flush_writers defaults to the smaller of (number of disks,
# number of cores), with a minimum of 2 and a maximum of 8.
#.
# If your data directories are backed by SSD, you should increase this
# to the number of cores.
#memtable_flush_writers: 8

# A fixed memory pool size in MB for for SSTable index summaries. If left
# empty, this will default to 5% of the heap size. If the memory usage of
# all index summaries exceeds this limit, SSTables with low read rates will
# shrink their index summaries in order to meet this limit.  However, this
# is a best-effort process. In extreme conditions Cassandra may need to use
# more than this amount of memory.
index_summary_capacity_in_mb:

# How frequently index summaries should be resampled.  This is done
# periodically to redistribute memory from the fixed-size pool to sstables
# proportional their recent read rates.  Setting to -1 will disable this
# process, leaving existing index summaries at their current sampling level.
index_summary_resize_interval_in_minutes: 60

# Whether to, when doing sequential writing, fsync() at intervals in
# order to force the operating system to flush the dirty
# buffers. Enable this to avoid sudden dirty buffer flushing from
# impacting read latencies. Almost always a good idea on SSDs; not
# necessarily on platters.
trickle_fsync: false
trickle_fsync_interval_in_kb: 10240

# TCP port, for commands and data
# For security reasons, you should not expose this port to the internet.  Firewall it if needed.
storage_port: 7000

# SSL port, for encrypted communication.  Unused unless enabled in
# encryption_options
# For security reasons, you should not expose this port to the internet.  Firewall it if needed.
ssl_storage_port: 7001

# Address or interface to bind to and tell other Cassandra nodes to connect to.
# You _must_ change this if you want multiple nodes to be able to communicate!
#
# Set listen_address OR listen_interface, not both. Interfaces must correspond
# to a single address, IP aliasing is not supported.
#
# Leaving it blank leaves it up to InetAddress.getLocalHost(). This
# will always do the Right Thing _if_ the node is properly configured
# (hostname, name resolution, etc), and the Right Thing is to use the
# address associated with the hostname (it might not be).
#
# Setting listen_address to 0.0.0.0 is always wrong.
#
# If you choose to specify the interface by name and the interface has an ipv4 and an ipv6 address
# you can specify which should be chosen using listen_interface_prefer_ipv6. If false the first ipv4
# address will be used. If true the first ipv6 address will be used. Defaults to false preferring
```

```
# ipv4. If there is only one address it will be selected regardless of ipv4/ipv6.
listen_address: 10.90.54.151
# listen_interface: eth0
# listen_interface_prefer_ipv6: false

# Address to broadcast to other Cassandra nodes
# Leaving this blank will set it to the same value as listen_address
# broadcast_address: 1.2.3.4

# Internode authentication backend, implementing IInternodeAuthenticator;
# used to allow/disallow connections from peer nodes.
# internode_authenticator: org.apache.cassandra.auth.AllowAllInternodeAuthenticator

# Whether to start the native transport server.
# Please note that the address on which the native transport is bound is the
# same as the rpc_address. The port however is different and specified below.
start_native_transport: true
# port for the CQL native transport to listen for clients on
# For security reasons, you should not expose this port to the internet.  Firewall it if needed.
native_transport_port: 9042
# The maximum threads for handling requests when the native transport is used.
# This is similar to rpc_max_threads though the default differs slightly (and
# there is no native_transport_min_threads, idle threads will always be stopped
# after 30 seconds).
# native_transport_max_threads: 128
#
# The maximum size of allowed frame. Frame (requests) larger than this will
# be rejected as invalid. The default is 256MB.
# native_transport_max_frame_size_in_mb: 256

# The maximum number of concurrent client connections.
# The default is -1, which means unlimited.
# native_transport_max_concurrent_connections: -1

# The maximum number of concurrent client connections per source ip.
# The default is -1, which means unlimited.
# native_transport_max_concurrent_connections_per_ip: -1


# Whether to start the thrift rpc server.
start_rpc: true

# The address or interface to bind the Thrift RPC service and native transport
# server to.
#
# Set rpc_address OR rpc_interface, not both. Interfaces must correspond
# to a single address, IP aliasing is not supported.
#
# Leaving rpc_address blank has the same effect as on listen_address
# (i.e. it will be based on the configured hostname of the node).
#
# Note that unlike listen_address, you can specify 0.0.0.0, but you must also
# set broadcast_rpc_address to a value other than 0.0.0.0.
#
# For security reasons, you should not expose this port to the internet.  Firewall it if needed.
#
# If you choose to specify the interface by name and the interface has an ipv4 and an ipv6 address
# you can specify which should be chosen using rpc_interface_prefer_ipv6. If false the first ipv4
# address will be used. If true the first ipv6 address will be used. Defaults to false preferring
# ipv4. If there is only one address it will be selected regardless of ipv4/ipv6.
rpc_address: 10.90.54.151
# rpc_interface: eth1
# rpc_interface_prefer_ipv6: false

# port for Thrift to listen for clients on
rpc_port: 9160

# RPC address to broadcast to drivers and other Cassandra nodes. This cannot
# be set to 0.0.0.0. If left blank, this will be set to the value of
# rpc_address. If rpc_address is set to 0.0.0.0, broadcast_rpc_address must
# be set.
# broadcast_rpc_address: 1.2.3.4
```

```
# enable or disable keepalive on rpc/native connections
rpc_keepalive: true

# Cassandra provides two out-of-the-box options for the RPC Server:
#
# sync  -> One thread per thrift connection. For a very large number of clients, memory
#          will be your limiting factor. On a 64 bit JVM, 180KB is the minimum stack size
#          per thread, and that will correspond to your use of virtual memory (but physical memory
#          may be limited depending on use of stack space).
#
# hsha  -> Stands for "half synchronous, half asynchronous." All thrift clients are handled
#          asynchronously using a small number of threads that does not vary with the amount
#          of thrift clients (and thus scales well to many clients). The rpc requests are still
#          synchronous (one thread per active request). If hsha is selected then it is essential
#          that rpc_max_threads is changed from the default value of unlimited.
#
# The default is sync because on Windows hsha is about 30% slower.  On Linux,
# sync/hsha performance is about the same, with hsha of course using less memory.
#
# Alternatively,  can provide your own RPC server by providing the fully-qualified class name
# of an o.a.c.t.TServerFactory that can create an instance of it.
rpc_server_type: sync

# Uncomment rpc_min|max_thread to set request pool size limits.
#
# Regardless of your choice of RPC server (see above), the number of maximum requests in the
# RPC thread pool dictates how many concurrent requests are possible (but if you are using the sync
# RPC server, it also dictates the number of clients that can be connected at all).
#
# The default is unlimited and thus provides no protection against clients overwhelming the server. You are
# encouraged to set a maximum that makes sense for you in production, but do keep in mind that
# rpc_max_threads represents the maximum number of client requests this server may execute concurrently.
#
# rpc_min_threads: 16
# rpc_max_threads: 2048

# uncomment to set socket buffer sizes on rpc connections
# rpc_send_buff_size_in_bytes:
# rpc_recv_buff_size_in_bytes:

# Uncomment to set socket buffer size for internode communication
# Note that when setting this, the buffer size is limited by net.core.wmem_max
# and when not setting it it is defined by net.ipv4.tcp_wmem
# See:
# /proc/sys/net/core/wmem_max
# /proc/sys/net/core/rmem_max
# /proc/sys/net/ipv4/tcp_wmem
# /proc/sys/net/ipv4/tcp_wmem
# and: man tcp
# internode_send_buff_size_in_bytes:
# internode_recv_buff_size_in_bytes:

# Frame size for thrift (maximum message length).
thrift_framed_transport_size_in_mb: 15

# Set to true to have Cassandra create a hard link to each sstable
# flushed or streamed locally in a backups/ subdirectory of the
# keyspace data.  Removing these links is the operator's
# responsibility.
incremental_backups: false

# Whether or not to take a snapshot before each compaction.  Be
# careful using this option, since Cassandra won't clean up the
# snapshots for you.  Mostly useful if you're paranoid when there
# is a data format change.
snapshot_before_compaction: false

# Whether or not a snapshot is taken of the data before keyspace truncation
# or dropping of column families. The STRONGLY advised default of true.
# should be used to provide data safety. If you set this flag to false, you will
# lose data on truncation or drop.
auto_snapshot: true
```

```
# When executing a scan, within or across a partition, we need to keep the
# tombstones seen in memory so we can return them to the coordinator, which
# will use them to make sure other replicas also know about the deleted rows.
# With workloads that generate a lot of tombstones, this can cause performance
# problems and even exaust the server heap.
# (http://www.datastax.com/dev/blog/cassandra-anti-patterns-queues-and-queue-like-datasets)
# Adjust the thresholds here if you understand the dangers and want to
# scan more tombstones anyway.  These thresholds may also be adjusted at runtime
# using the StorageService mbean.
tombstone_warn_threshold: 1000
tombstone_failure_threshold: 100000

# Granularity of the collation index of rows within a partition.
# Increase if your rows are large, or if you have a very large
# number of rows per partition.  The competing goals are these:
#   1) a smaller granularity means more index entries are generated
#      and looking up rows withing the partition by collation column
#      is faster
#   2) but, Cassandra will keep the collation index in memory for hot
#      rows (as part of the key cache), so a larger granularity means
#      you can cache more hot rows
column_index_size_in_kb: 64


# Log WARN on any batch size exceeding this value. 5kb per batch by default.
# Caution should be taken on increasing the size of this threshold as it can lead to node instability.
batch_size_warn_threshold_in_kb: 5

# Number of simultaneous compactions to allow, NOT including
# validation "compactions" for anti-entropy repair.  Simultaneous
# compactions can help preserve read performance in a mixed read/write
# workload, by mitigating the tendency of small sstables to accumulate
# during a single long running compactions. The default is usually
# fine and if you experience problems with compaction running too
# slowly or too fast, you should look at
# compaction_throughput_mb_per_sec first.
#
# concurrent_compactors defaults to the smaller of (number of disks,
# number of cores), with a minimum of 2 and a maximum of 8.
#.
# If your data directories are backed by SSD, you should increase this
# to the number of cores.
#concurrent_compactors: 1

# Throttles compaction to the given total throughput across the entire
# system. The faster you insert data, the faster you need to compact in
# order to keep the sstable count down, but in general, setting this to
# 16 to 32 times the rate you are inserting data is more than sufficient.
# Setting this to 0 disables throttling. Note that this account for all types
# of compaction, including validation compaction.
compaction_throughput_mb_per_sec: 16

# Log a warning when compacting partitions larger than this value
compaction_large_partition_warning_threshold_mb: 100

# When compacting, the replacement sstable(s) can be opened before they
# are completely written, and used in place of the prior sstables for
# any range that has been written. This helps to smoothly transfer reads.
# between the sstables, reducing page cache churn and keeping hot rows hot
sstable_preemptive_open_interval_in_mb: 50

# Throttles all outbound streaming file transfers on this node to the
# given total throughput in Mbps. This is necessary because Cassandra does
# mostly sequential IO when streaming data during bootstrap or repair, which
# can lead to saturating the network connection and degrading rpc performance.
# When unset, the default is 200 Mbps or 25 MB/s.
# stream_throughput_outbound_megabits_per_sec: 200

# Throttles all streaming file transfer between the datacenters,
# this setting allows users to throttle inter dc stream throughput in addition
# to throttling all network stream traffic as configured with
# stream_throughput_outbound_megabits_per_sec
# inter_dc_stream_throughput_outbound_megabits_per_sec:
```

```
# How long the coordinator should wait for read operations to complete
read_request_timeout_in_ms: 5000
# How long the coordinator should wait for seq or index scans to complete
range_request_timeout_in_ms: 10000
# How long the coordinator should wait for writes to complete
write_request_timeout_in_ms: 2000
# How long the coordinator should wait for counter writes to complete
counter_write_request_timeout_in_ms: 5000
# How long a coordinator should continue to retry a CAS operation
# that contends with other proposals for the same row
cas_contention_timeout_in_ms: 1000
# How long the coordinator should wait for truncates to complete
# (This can be much longer, because unless auto_snapshot is disabled
# we need to flush first so we can snapshot before removing the data.)
truncate_request_timeout_in_ms: 60000
# The default timeout for other, miscellaneous operations
request_timeout_in_ms: 10000

# Enable operation timeout information exchange between nodes to accurately
# measure request timeouts.  If disabled, replicas will assume that requests
# were forwarded to them instantly by the coordinator, which means that
# under overload conditions we will waste that much extra time processing.
# already-timed-out requests.
#
# Warning: before enabling this property make sure to ntp is installed
# and the times are synchronized between the nodes.
cross_node_timeout: false

# Enable socket timeout for streaming operation.
# When a timeout occurs during streaming, streaming is retried from the start
# of the current file. This _can_ involve re-streaming an important amount of
# data, so you should avoid setting the value too low.
# Default value is 3600000, which means streams timeout after an hour.
# streaming_socket_timeout_in_ms: 3600000

# phi value that must be reached for a host to be marked down.
# most users should never need to adjust this.
# phi_convict_threshold: 8

# endpoint_snitch -- Set this to a class that implements
# IEndpointSnitch.  The snitch has two functions:
# - it teaches Cassandra enough about your network topology to route
#   requests efficiently
# - it allows Cassandra to spread replicas around your cluster to avoid
#   correlated failures. It does this by grouping machines into
#   "datacenters" and "racks."  Cassandra will do its best not to have
#   more than one replica on the same "rack" (which may not actually
#   be a physical location)
#
# IF YOU CHANGE THE SNITCH AFTER DATA IS INSERTED INTO THE CLUSTER,
# YOU MUST RUN A FULL REPAIR, SINCE THE SNITCH AFFECTS WHERE REPLICAS
# ARE PLACED.
#
# Out of the box, Cassandra provides
#  - SimpleSnitch:
#    Treats Strategy order as proximity. This can improve cache
#    locality when disabling read repair.  Only appropriate for
#    single-datacenter deployments.
#  - GossipingPropertyFileSnitch
#    This should be your go-to snitch for production use.  The rack
#    and datacenter for the local node are defined in
#    cassandra-rackdc.properties and propagated to other nodes via
#    gossip.  If cassandra-topology.properties exists, it is used as a
#    fallback, allowing migration from the PropertyFileSnitch.
#  - PropertyFileSnitch:
#    Proximity is determined by rack and data center, which are
#    explicitly configured in cassandra-topology.properties.
#  - Ec2Snitch:
#    Appropriate for EC2 deployments in a single Region. Loads Region
#    and Availability Zone information from the EC2 API. The Region is
#    treated as the datacenter, and the Availability Zone as the rack.
#    Only private IPs are used, so this will not work across multiple
```

```
#    Regions.
#  - Ec2MultiRegionSnitch:
#    Uses public IPs as broadcast_address to allow cross-region
#    connectivity.  (Thus, you should set seed addresses to the public
#    IP as well.) You will need to open the storage_port or
#    ssl_storage_port on the public IP firewall.  (For intra-Region
#    traffic, Cassandra will switch to the private IP after
#    establishing a connection.)
#  - RackInferringSnitch:
#    Proximity is determined by rack and data center, which are
#    assumed to correspond to the 3rd and 2nd octet of each node's IP
#    address, respectively.  Unless this happens to match your
#    deployment conventions, this is best used as an example of
#    writing a custom Snitch class and is provided in that spirit.
#
# You can use a custom Snitch by setting this to the full class name
# of the snitch, which will be assumed to be on your classpath.
endpoint_snitch: GossipingPropertyFileSnitch

# controls how often to perform the more expensive part of host score
# calculation
dynamic_snitch_update_interval_in_ms: 100.
# controls how often to reset all host scores, allowing a bad host to
# possibly recover
dynamic_snitch_reset_interval_in_ms: 600000
# if set greater than zero and read_repair_chance is < 1.0, this will allow
# 'pinning' of replicas to hosts in order to increase cache capacity.
# The badness threshold will control how much worse the pinned host has to be
# before the dynamic snitch will prefer other replicas over it.  This is
# expressed as a double which represents a percentage.  Thus, a value of
# 0.2 means Cassandra would continue to prefer the static snitch values
# until the pinned host was 20% worse than the fastest.
dynamic_snitch_badness_threshold: 0.1

# request_scheduler -- Set this to a class that implements
# RequestScheduler, which will schedule incoming client requests
# according to the specific policy. This is useful for multi-tenancy
# with a single Cassandra cluster.
# NOTE: This is specifically for requests from the client and does
# not affect inter node communication.
# org.apache.cassandra.scheduler.NoScheduler - No scheduling takes place
# org.apache.cassandra.scheduler.RoundRobinScheduler - Round robin of
# client requests to a node with a separate queue for each
# request_scheduler_id. The scheduler is further customized by
# request_scheduler_options as described below.
request_scheduler: org.apache.cassandra.scheduler.NoScheduler

# Scheduler Options vary based on the type of scheduler
# NoScheduler - Has no options
# RoundRobin
#  - throttle_limit -- The throttle_limit is the number of in-flight
#                      requests per client.  Requests beyond.
#                      that limit are queued up until
#                      running requests can complete.
#                      The value of 80 here is twice the number of
#                      concurrent_reads + concurrent_writes.
#  - default_weight -- default_weight is optional and allows for
#                      overriding the default which is 1.
#  - weights -- Weights are optional and will default to 1 or the
#               overridden default_weight. The weight translates into how
#               many requests are handled during each turn of the
#               RoundRobin, based on the scheduler id.
#
# request_scheduler_options:
#    throttle_limit: 80
#    default_weight: 5
#    weights:
#      Keyspace1: 1
#      Keyspace2: 5

# request_scheduler_id -- An identifier based on which to perform
# the request scheduling. Currently the only valid option is keyspace.
# request_scheduler_id: keyspace
```

```
# Enable or disable inter-node encryption
# Default settings are TLS v1, RSA 1024-bit keys (it is imperative that
# users generate their own keys) TLS_RSA_WITH_AES_128_CBC_SHA as the cipher
# suite for authentication, key exchange and encryption of the actual data transfers.
# Use the DHE/ECDHE ciphers if running in FIPS 140 compliant mode.
# NOTE: No custom encryption options are enabled at the moment
# The available internode options are : all, none, dc, rack
#
# If set to dc cassandra will encrypt the traffic between the DCs
# If set to rack cassandra will encrypt the traffic between the racks
#
# The passwords used in these options must match the passwords used when generating
# the keystore and truststore.  For instructions on generating these files, see:
# http://download.oracle.com/javase/6/docs/technotes/guides/security/jsse/JSSERefGuide.html#CreateKeystore
#
server_encryption_options:
    internode_encryption: none
    keystore: conf/.keystore
    keystore_password: cassandra
    truststore: conf/.truststore
    truststore_password: cassandra
    # More advanced defaults below:
    # protocol: TLS
    # algorithm: SunX509
    # store_type: JKS
    # cipher_suites:
[TLS_RSA_WITH_AES_128_CBC_SHA,TLS_RSA_WITH_AES_256_CBC_SHA,TLS_DHE_RSA_WITH_AES_
128_CBC_SHA,TLS_DHE_RSA_WITH_AES_256_CBC_SHA,TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,
TLS_ECDHE_RSA_W
    # require_client_auth: false

# enable or disable client/server encryption.
client_encryption_options:
    enabled: false
    keystore: conf/.keystore
    keystore_password: cassandra
    # require_client_auth: false
    # Set trustore and truststore_password if require_client_auth is true
    # truststore: conf/.truststore
    # truststore_password: cassandra
    # More advanced defaults below:
    # protocol: TLS
    # algorithm: SunX509
    # store_type: JKS
    # cipher_suites:
[TLS_RSA_WITH_AES_128_CBC_SHA,TLS_RSA_WITH_AES_256_CBC_SHA,TLS_DHE_RSA_WITH_AES_
128_CBC_SHA,TLS_DHE_RSA_WITH_AES_256_CBC_SHA,TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,
TLS_ECDHE_RSA_W

# internode_compression controls whether traffic between nodes is
# compressed.
# can be:  all  - all traffic is compressed
#          dc   - traffic between different datacenters is compressed
#          none - nothing is compressed.
internode_compression: all

# Enable or disable tcp_nodelay for inter-dc communication.
# Disabling it will result in larger (but fewer) network packets being sent,
# reducing overhead from the TCP protocol itself, at the cost of increasing
# latency if you block for cross-datacenter responses.
inter_dc_tcp_nodelay: false

# GC Pauses greater than gc_warn_threshold_in_ms will be logged at WARN level
# Adjust the threshold based on your application throughput requirement
# By default, Cassandra logs GC Pauses greater than 200 ms at INFO level
# gc_warn_threshold_in_ms: 1000
#
#
#auto_bootstrap: false
```

I. Nazwa serwera: cassandra2

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	cassandra2
IP / Domena	10.90.54.152
OS	CentOS 7
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	100 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Bazodanowy

2. Konfiguracja .hosts

```
127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6
```

3. Interfejsy sieciowe

eth0 10.90.54.152

4. Zainstalowane usługi

a. Cassandra

Parametr	Wartość
Lokalizacja instalacji	/etc/cassandra
Zawartość pliku konfiguracyjnego #1	/etc/cassandra/conf # Cassandra storage config YAML. # NOTE: # See http://wiki.apache.org/cassandra/StorageConfiguration for # full explanations of configuration directives # /NOTE # The name of the cluster. This is mainly used to prevent machines in # one logical cluster from joining another. cluster_name: 'EpoCluster' # This defines the number of tokens randomly assigned to this node on the ring # The more tokens, relative to other nodes, the larger the proportion of data # that this node will store. You probably want all nodes to have the same number # of tokens assuming they have equal hardware capability. # # If you leave this unspecified, Cassandra will use the default of 1 token for legacy compatibility, # and will use the initial_token as described below. # # Specifying initial_token will override this setting on the node's initial start, # on subsequent starts, this setting will apply even if initial token is set. # # If you already have a cluster with 1 token per node, and wish to migrate to.

```
# multiple tokens per node, see http://wiki.apache.org/cassandra/Operations
num_tokens: 256

# initial_token allows you to specify tokens manually.  While you can use # it with
# vnodes (num_tokens > 1, above) -- in which case you should provide a.
# comma-separated list -- it's primarily used when adding nodes # to legacy clusters.
# that do not have vnodes enabled.
# initial_token:

# See http://wiki.apache.org/cassandra/HintedHandoff
# May either be "true" or "false" to enable globally, or contain a list
# of data centers to enable per-datacenter.
# hinted_handoff_enabled: DC1,DC2
hinted_handoff_enabled: true
# this defines the maximum amount of time a dead host will have hints
# generated.  After it has been dead this long, new hints for it will not be
# created until it has been seen alive and gone down again.
max_hint_window_in_ms: 10800000 # 3 hours
# Maximum throttle in KBs per second, per delivery thread.  This will be
# reduced proportionally to the number of nodes in the cluster.  (If there
# are two nodes in the cluster, each delivery thread will use the maximum
# rate; if there are three, each will throttle to half of the maximum,
# since we expect two nodes to be delivering hints simultaneously.)
hinted_handoff_throttle_in_kb: 1024
# Number of threads with which to deliver hints;
# Consider increasing this number when you have multi-dc deployments, since
# cross-dc handoff tends to be slower
max_hints_delivery_threads: 2

# Maximum throttle in KBs per second, total. This will be
# reduced proportionally to the number of nodes in the cluster.
batchlog_replay_throttle_in_kb: 1024

# Authentication backend, implementing IAuthenticator; used to identify users
# Out of the box, Cassandra provides org.apache.cassandra.auth.{AllowAllAuthenticator,
# PasswordAuthenticator}.
#
# - AllowAllAuthenticator performs no checks - set it to disable authentication.
# - PasswordAuthenticator relies on username/password pairs to authenticate
#   users. It keeps usernames and hashed passwords in system_auth.credentials table.
#   Please increase system_auth keyspace replication factor if you use this authenticator.
authenticator: PasswordAuthenticator

# Authorization backend, implementing IAuthorizer; used to limit access/provide permissions
# Out of the box, Cassandra provides org.apache.cassandra.auth.{AllowAllAuthorizer,
# CassandraAuthorizer}.
#
# - AllowAllAuthorizer allows any action to any user - set it to disable authorization.
# - CassandraAuthorizer stores permissions in system_auth.permissions table. Please
#   increase system_auth keyspace replication factor if you use this authorizer.
authorizer: CassandraAuthorizer

# Validity period for permissions cache (fetching permissions can be an
# expensive operation depending on the authorizer, CassandraAuthorizer is
# one example). Defaults to 2000, set to 0 to disable.
# Will be disabled automatically for AllowAllAuthorizer.
permissions_validity_in_ms: 2000

# Refresh interval for permissions cache (if enabled).
# After this interval, cache entries become eligible for refresh. Upon next
# access, an async reload is scheduled and the old value returned until it
# completes. If permissions_validity_in_ms is non-zero, then this must be
# also.
# Defaults to the same value as permissions_validity_in_ms.
# permissions_update_interval_in_ms: 1000

# The partitioner is responsible for distributing groups of rows (by
# partition key) across nodes in the cluster.  You should leave this
# alone for new clusters.  The partitioner can NOT be changed without
# reloading all data, so when upgrading you should set this to the
# same partitioner you were already using.
#
# Besides Murmur3Partitioner, partitioners included for backwards
```

```
# compatibility include RandomPartitioner, ByteOrderedPartitioner, and
# OrderPreservingPartitioner.
#
partitioner: org.apache.cassandra.dht.Murmur3Partitioner

# Directories where Cassandra should store data on disk.  Cassandra
# will spread data evenly across them, subject to the granularity of
# the configured compaction strategy.
# If not set, the default directory is $CASSANDRA_HOME/data/data.
data_file_directories:
    - /opt/cassandra/data

# commit log.  when running on magnetic HDD, this should be a
# separate spindle than the data directories.
# If not set, the default directory is $CASSANDRA_HOME/data/commitlog.
commitlog_directory: /opt/cassandra/commitlog

# policy for data disk failures:
# die: shut down gossip and client transports and kill the JVM for any fs errors or
#      single-sstable errors, so the node can be replaced.
# stop_paranoid: shut down gossip and client transports even for single-sstable errors,
#               kill the JVM for errors during startup.
# stop: shut down gossip and client transports, leaving the node effectively dead, but
#       can still be inspected via JMX, kill the JVM for errors during startup.
# best_effort: stop using the failed disk and respond to requests based on
#              remaining available sstables.  This means you WILL see obsolete
#              data at CL.ONE!
# ignore: ignore fatal errors and let requests fail, as in pre-1.2 Cassandra
disk_failure_policy: stop

# policy for commit disk failures:
# die: shut down gossip and Thrift and kill the JVM, so the node can be replaced.
# stop: shut down gossip and Thrift, leaving the node effectively dead, but
#       can still be inspected via JMX.
# stop_commit: shutdown the commit log, letting writes collect but
#              continuing to service reads, as in pre-2.0.5 Cassandra
# ignore: ignore fatal errors and let the batches fail
commit_failure_policy: stop

# Maximum size of the key cache in memory.
#
# Each key cache hit saves 1 seek and each row cache hit saves 2 seeks at the
# minimum, sometimes more. The key cache is fairly tiny for the amount of
# time it saves, so it's worthwhile to use it at large numbers.
# The row cache saves even more time, but must contain the entire row,
# so it is extremely space-intensive. It's best to only use the
# row cache if you have hot rows or static rows.
#
# NOTE: if you reduce the size, you may not get you hottest keys loaded on startup.
#
# Default value is empty to make it "auto" (min(5% of Heap (in MB), 100MB)). Set to 0 to disable key cache.
key_cache_size_in_mb:

# Duration in seconds after which Cassandra should
# save the key cache. Caches are saved to saved_caches_directory as
# specified in this configuration file.
#
# Saved caches greatly improve cold-start speeds, and is relatively cheap in
# terms of I/O for the key cache. Row cache saving is much more expensive and
# has limited use.
#
# Default is 14400 or 4 hours.
key_cache_save_period: 14400

# Number of keys from the key cache to save
# Disabled by default, meaning all keys are going to be saved
# key_cache_keys_to_save: 100

# Maximum size of the row cache in memory.
# NOTE: if you reduce the size, you may not get you hottest keys loaded on startup.
#
# Default value is 0, to disable row caching.
row_cache_size_in_mb: 0
```

```
# Duration in seconds after which Cassandra should
# save the row cache. Caches are saved to saved_caches_directory as specified
# in this configuration file.
#
# Saved caches greatly improve cold-start speeds, and is relatively cheap in
# terms of I/O for the key cache. Row cache saving is much more expensive and
# has limited use.
#
# Default is 0 to disable saving the row cache.
row_cache_save_period: 0

# Number of keys from the row cache to save
# Disabled by default, meaning all keys are going to be saved
# row_cache_keys_to_save: 100

# Maximum size of the counter cache in memory.
#
# Counter cache helps to reduce counter locks' contention for hot counter cells.
# In case of RF = 1 a counter cache hit will cause Cassandra to skip the read before
# write entirely. With RF > 1 a counter cache hit will still help to reduce the duration
# of the lock hold, helping with hot counter cell updates, but will not allow skipping
# the read entirely. Only the local (clock, count) tuple of a counter cell is kept
# in memory, not the whole counter, so it's relatively cheap.
#
# NOTE: if you reduce the size, you may not get you hottest keys loaded on startup.
#
# Default value is empty to make it "auto" (min(2.5% of Heap (in MB), 50MB)). Set to 0 to disable counter cache.
# NOTE: if you perform counter deletes and rely on low gcgs, you should disable the counter cache.
counter_cache_size_in_mb:

# Duration in seconds after which Cassandra should
# save the counter cache (keys only). Caches are saved to saved_caches_directory as
# specified in this configuration file.
#
# Default is 7200 or 2 hours.
counter_cache_save_period: 7200

# Number of keys from the counter cache to save
# Disabled by default, meaning all keys are going to be saved
# counter_cache_keys_to_save: 100

# The off-heap memory allocator.  Affects storage engine metadata as
# well as caches.  Experiments show that JEMAlloc saves some memory
# than the native GCC allocator (i.e., JEMalloc is more
# fragmentation-resistant).
#.
# Supported values are: NativeAllocator, JEMallocAllocator
#
# If you intend to use JEMallocAllocator you have to install JEMalloc as library and
# modify cassandra-env.sh as directed in the file.
#
# Defaults to NativeAllocator
# memory_allocator: NativeAllocator

# saved caches
# If not set, the default directory is $CASSANDRA_HOME/data/saved_caches.
saved_caches_directory: /opt/cassandra/saved_caches

# commitlog_sync may be either "periodic" or "batch.".
#.
# When in batch mode, Cassandra won't ack writes until the commit log
# has been fsynced to disk.  It will wait
# commitlog_sync_batch_window_in_ms milliseconds between fsyncs.
# This window should be kept short because the writer threads will
# be unable to do extra work while waiting.  (You may need to increase
# concurrent_writes for the same reason.)
#
# commitlog_sync: batch
# commitlog_sync_batch_window_in_ms: 2
#
# the other option is "periodic" where writes may be acked immediately
# and the CommitLog is simply synced every commitlog_sync_period_in_ms
```

```
# milliseconds..
commitlog_sync: periodic
commitlog_sync_period_in_ms: 10000

# The size of the individual commitlog file segments.  A commitlog
# segment may be archived, deleted, or recycled once all the data
# in it (potentially from each columnfamily in the system) has been
# flushed to sstables...
#
# The default size is 32, which is almost always fine, but if you are
# archiving commitlog segments (see commitlog_archiving.properties),
# then you probably want a finer granularity of archiving; 8 or 16 MB
# is reasonable.
commitlog_segment_size_in_mb: 32

# Reuse commit log files when possible. The default is false, and this
# feature will be removed entirely in future versions of Cassandra.
#commitlog_segment_recycling: false

# any class that implements the SeedProvider interface and has a
# constructor that takes a Map<String, String> of parameters will do.
seed_provider:
   # Addresses of hosts that are deemed contact points..
   # Cassandra nodes use this list of hosts to find each other and learn
   # the topology of the ring.  You must change this if you are running
   # multiple nodes!
   - class_name: org.apache.cassandra.locator.SimpleSeedProvider
     parameters:
         # seeds is actually a comma-delimited list of addresses.
         # Ex: "<ip1>,<ip2>,<ip3>"
         - seeds: "10.90.54.151,10.90.54.152,10.90.54.153,10.90.54.154"

# For workloads with more data than can fit in memory, Cassandra's
# bottleneck will be reads that need to fetch data from
# disk. "concurrent_reads" should be set to (16 * number_of_drives) in
# order to allow the operations to enqueue low enough in the stack
# that the OS and drives can reorder them. Same applies to
# "concurrent_counter_writes", since counter writes read the current
# values before incrementing and writing them back.
#
# On the other hand, since writes are almost never IO bound, the ideal
# number of "concurrent_writes" is dependent on the number of cores in
# your system; (8 * number_of_cores) is a good rule of thumb.
concurrent_reads: 32
concurrent_writes: 32
concurrent_counter_writes: 32

# Total memory to use for sstable-reading buffers.  Defaults to
# the smaller of 1/4 of heap or 512MB.
# file_cache_size_in_mb: 512

# Total permitted memory to use for memtables. Cassandra will stop.
# accepting writes when the limit is exceeded until a flush completes,
# and will trigger a flush based on memtable_cleanup_threshold
# If omitted, Cassandra will set both to 1/4 the size of the heap.
# memtable_heap_space_in_mb: 2048
# memtable_offheap_space_in_mb: 2048

# Ratio of occupied non-flushing memtable size to total permitted size
# that will trigger a flush of the largest memtable.  Lager mct will
# mean larger flushes and hence less compaction, but also less concurrent
# flush activity which can make it difficult to keep your disks fed
# under heavy write load.
#
# memtable_cleanup_threshold defaults to 1 / (memtable_flush_writers + 1)
# memtable_cleanup_threshold: 0.11

# Specify the way Cassandra allocates and manages memtable memory.
# Options are:
#   heap_buffers:    on heap nio buffers
#   offheap_buffers: off heap (direct) nio buffers
#   offheap_objects: native memory, eliminating nio buffer heap overhead
memtable_allocation_type: heap_buffers
```

```
# Total space to use for commitlogs.  Since commitlog segments are
# mmapped, and hence use up address space, the default size is 32
# on 32-bit JVMs, and 8192 on 64-bit JVMs.
#
# If space gets above this value (it will round up to the next nearest
# segment multiple), Cassandra will flush every dirty CF in the oldest
# segment and remove it.  So a small total commitlog space will tend
# to cause more flush activity on less-active columnfamilies.
# commitlog_total_space_in_mb: 8192

# This sets the amount of memtable flush writer threads.  These will
# be blocked by disk io, and each one will hold a memtable in memory
# while blocked..
#
# memtable_flush_writers defaults to the smaller of (number of disks,
# number of cores), with a minimum of 2 and a maximum of 8.
#.
# If your data directories are backed by SSD, you should increase this
# to the number of cores.
#memtable_flush_writers: 8

# A fixed memory pool size in MB for for SSTable index summaries. If left
# empty, this will default to 5% of the heap size. If the memory usage of
# all index summaries exceeds this limit, SSTables with low read rates will
# shrink their index summaries in order to meet this limit.  However, this
# is a best-effort process. In extreme conditions Cassandra may need to use
# more than this amount of memory.
index_summary_capacity_in_mb:

# How frequently index summaries should be resampled.  This is done
# periodically to redistribute memory from the fixed-size pool to sstables
# proportional their recent read rates.  Setting to -1 will disable this
# process, leaving existing index summaries at their current sampling level.
index_summary_resize_interval_in_minutes: 60

# Whether to, when doing sequential writing, fsync() at intervals in
# order to force the operating system to flush the dirty
# buffers. Enable this to avoid sudden dirty buffer flushing from
# impacting read latencies. Almost always a good idea on SSDs; not
# necessarily on platters.
trickle_fsync: false
trickle_fsync_interval_in_kb: 10240

# TCP port, for commands and data
# For security reasons, you should not expose this port to the internet.  Firewall it if needed.
storage_port: 7000

# SSL port, for encrypted communication.  Unused unless enabled in
# encryption_options
# For security reasons, you should not expose this port to the internet.  Firewall it if needed.
ssl_storage_port: 7001

# Address or interface to bind to and tell other Cassandra nodes to connect to.
# You _must_ change this if you want multiple nodes to be able to communicate!
#
# Set listen_address OR listen_interface, not both. Interfaces must correspond
# to a single address, IP aliasing is not supported.
#
# Leaving it blank leaves it up to InetAddress.getLocalHost(). This
# will always do the Right Thing _if_ the node is properly configured
# (hostname, name resolution, etc), and the Right Thing is to use the
# address associated with the hostname (it might not be).
#
# Setting listen_address to 0.0.0.0 is always wrong.
#
# If you choose to specify the interface by name and the interface has an ipv4 and an ipv6 address
# you can specify which should be chosen using listen_interface_prefer_ipv6. If false the first ipv4
# address will be used. If true the first ipv6 address will be used. Defaults to false preferring
# ipv4. If there is only one address it will be selected regardless of ipv4/ipv6.
listen_address: 10.90.54.151
# listen_interface: eth0
# listen_interface_prefer_ipv6: false
```

```
# Address to broadcast to other Cassandra nodes
# Leaving this blank will set it to the same value as listen_address
# broadcast_address: 1.2.3.4

# Internode authentication backend, implementing IInternodeAuthenticator;
# used to allow/disallow connections from peer nodes.
# internode_authenticator: org.apache.cassandra.auth.AllowAllInternodeAuthenticator

# Whether to start the native transport server.
# Please note that the address on which the native transport is bound is the
# same as the rpc_address. The port however is different and specified below.
start_native_transport: true
# port for the CQL native transport to listen for clients on
# For security reasons, you should not expose this port to the internet.  Firewall it if needed.
native_transport_port: 9042
# The maximum threads for handling requests when the native transport is used.
# This is similar to rpc_max_threads though the default differs slightly (and
# there is no native_transport_min_threads, idle threads will always be stopped
# after 30 seconds).
# native_transport_max_threads: 128
#
# The maximum size of allowed frame. Frame (requests) larger than this will
# be rejected as invalid. The default is 256MB.
# native_transport_max_frame_size_in_mb: 256

# The maximum number of concurrent client connections.
# The default is -1, which means unlimited.
# native_transport_max_concurrent_connections: -1

# The maximum number of concurrent client connections per source ip.
# The default is -1, which means unlimited.
# native_transport_max_concurrent_connections_per_ip: -1


# Whether to start the thrift rpc server.
start_rpc: true

# The address or interface to bind the Thrift RPC service and native transport
# server to.
#
# Set rpc_address OR rpc_interface, not both. Interfaces must correspond
# to a single address, IP aliasing is not supported.
#
# Leaving rpc_address blank has the same effect as on listen_address
# (i.e. it will be based on the configured hostname of the node).
#
# Note that unlike listen_address, you can specify 0.0.0.0, but you must also
# set broadcast_rpc_address to a value other than 0.0.0.0.
#
# For security reasons, you should not expose this port to the internet.  Firewall it if needed.
#
# If you choose to specify the interface by name and the interface has an ipv4 and an ipv6 address
# you can specify which should be chosen using rpc_interface_prefer_ipv6. If false the first ipv4
# address will be used. If true the first ipv6 address will be used. Defaults to false preferring
# ipv4. If there is only one address it will be selected regardless of ipv4/ipv6.
rpc_address: 10.90.54.151
# rpc_interface: eth1
# rpc_interface_prefer_ipv6: false

# port for Thrift to listen for clients on
rpc_port: 9160

# RPC address to broadcast to drivers and other Cassandra nodes. This cannot
# be set to 0.0.0.0. If left blank, this will be set to the value of
# rpc_address. If rpc_address is set to 0.0.0.0, broadcast_rpc_address must
# be set.
# broadcast_rpc_address: 1.2.3.4

# enable or disable keepalive on rpc/native connections
rpc_keepalive: true

# Cassandra provides two out-of-the-box options for the RPC Server:
```

```
#
# sync  -> One thread per thrift connection. For a very large number of clients, memory
#          will be your limiting factor. On a 64 bit JVM, 180KB is the minimum stack size
#          per thread, and that will correspond to your use of virtual memory (but physical memory
#          may be limited depending on use of stack space).
#
# hsha  -> Stands for "half synchronous, half asynchronous." All thrift clients are handled
#          asynchronously using a small number of threads that does not vary with the amount
#          of thrift clients (and thus scales well to many clients). The rpc requests are still
#          synchronous (one thread per active request). If hsha is selected then it is essential
#          that rpc_max_threads is changed from the default value of unlimited.
#
# The default is sync because on Windows hsha is about 30% slower.  On Linux,
# sync/hsha performance is about the same, with hsha of course using less memory.
#
# Alternatively,  can provide your own RPC server by providing the fully-qualified class name
# of an o.a.c.t.TServerFactory that can create an instance of it.
rpc_server_type: sync

# Uncomment rpc_min|max_thread to set request pool size limits.
#
# Regardless of your choice of RPC server (see above), the number of maximum requests in the
# RPC thread pool dictates how many concurrent requests are possible (but if you are using the sync
# RPC server, it also dictates the number of clients that can be connected at all).
#
# The default is unlimited and thus provides no protection against clients overwhelming the server. You are
# encouraged to set a maximum that makes sense for you in production, but do keep in mind that
# rpc_max_threads represents the maximum number of client requests this server may execute concurrently.
#
# rpc_min_threads: 16
# rpc_max_threads: 2048

# uncomment to set socket buffer sizes on rpc connections
# rpc_send_buff_size_in_bytes:
# rpc_recv_buff_size_in_bytes:

# Uncomment to set socket buffer size for internode communication
# Note that when setting this, the buffer size is limited by net.core.wmem_max
# and when not setting it it is defined by net.ipv4.tcp_wmem
# See:
# /proc/sys/net/core/wmem_max
# /proc/sys/net/core/rmem_max
# /proc/sys/net/ipv4/tcp_wmem
# /proc/sys/net/ipv4/tcp_wmem
# and: man tcp
# internode_send_buff_size_in_bytes:
# internode_recv_buff_size_in_bytes:

# Frame size for thrift (maximum message length).
thrift_framed_transport_size_in_mb: 15

# Set to true to have Cassandra create a hard link to each sstable
# flushed or streamed locally in a backups/ subdirectory of the
# keyspace data.  Removing these links is the operator's
# responsibility.
incremental_backups: false

# Whether or not to take a snapshot before each compaction.  Be
# careful using this option, since Cassandra won't clean up the
# snapshots for you.  Mostly useful if you're paranoid when there
# is a data format change.
snapshot_before_compaction: false

# Whether or not a snapshot is taken of the data before keyspace truncation
# or dropping of column families. The STRONGLY advised default of true.
# should be used to provide data safety. If you set this flag to false, you will
# lose data on truncation or drop.
auto_snapshot: true

# When executing a scan, within or across a partition, we need to keep the
# tombstones seen in memory so we can return them to the coordinator, which
# will use them to make sure other replicas also know about the deleted rows.
# With workloads that generate a lot of tombstones, this can cause performance
```

```
# problems and even exaust the server heap.
# (http://www.datastax.com/dev/blog/cassandra-anti-patterns-queues-and-queue-like-datasets)
# Adjust the thresholds here if you understand the dangers and want to
# scan more tombstones anyway.  These thresholds may also be adjusted at runtime
# using the StorageService mbean.
tombstone_warn_threshold: 1000
tombstone_failure_threshold: 100000

# Granularity of the collation index of rows within a partition.
# Increase if your rows are large, or if you have a very large
# number of rows per partition.  The competing goals are these:
#   1) a smaller granularity means more index entries are generated
#      and looking up rows withing the partition by collation column
#      is faster
#   2) but, Cassandra will keep the collation index in memory for hot
#      rows (as part of the key cache), so a larger granularity means
#      you can cache more hot rows
column_index_size_in_kb: 64


# Log WARN on any batch size exceeding this value. 5kb per batch by default.
# Caution should be taken on increasing the size of this threshold as it can lead to node instability.
batch_size_warn_threshold_in_kb: 5

# Number of simultaneous compactions to allow, NOT including
# validation "compactions" for anti-entropy repair.  Simultaneous
# compactions can help preserve read performance in a mixed read/write
# workload, by mitigating the tendency of small sstables to accumulate
# during a single long running compactions. The default is usually
# fine and if you experience problems with compaction running too
# slowly or too fast, you should look at
# compaction_throughput_mb_per_sec first.
#
# concurrent_compactors defaults to the smaller of (number of disks,
# number of cores), with a minimum of 2 and a maximum of 8.
#.
# If your data directories are backed by SSD, you should increase this
# to the number of cores.
#concurrent_compactors: 1

# Throttles compaction to the given total throughput across the entire
# system. The faster you insert data, the faster you need to compact in
# order to keep the sstable count down, but in general, setting this to
# 16 to 32 times the rate you are inserting data is more than sufficient.
# Setting this to 0 disables throttling. Note that this account for all types
# of compaction, including validation compaction.
compaction_throughput_mb_per_sec: 16

# Log a warning when compacting partitions larger than this value
compaction_large_partition_warning_threshold_mb: 100

# When compacting, the replacement sstable(s) can be opened before they
# are completely written, and used in place of the prior sstables for
# any range that has been written. This helps to smoothly transfer reads.
# between the sstables, reducing page cache churn and keeping hot rows hot
sstable_preemptive_open_interval_in_mb: 50

# Throttles all outbound streaming file transfers on this node to the
# given total throughput in Mbps. This is necessary because Cassandra does
# mostly sequential IO when streaming data during bootstrap or repair, which
# can lead to saturating the network connection and degrading rpc performance.
# When unset, the default is 200 Mbps or 25 MB/s.
# stream_throughput_outbound_megabits_per_sec: 200

# Throttles all streaming file transfer between the datacenters,
# this setting allows users to throttle inter dc stream throughput in addition
# to throttling all network stream traffic as configured with
# stream_throughput_outbound_megabits_per_sec
# inter_dc_stream_throughput_outbound_megabits_per_sec:

# How long the coordinator should wait for read operations to complete
read_request_timeout_in_ms: 5000
# How long the coordinator should wait for seq or index scans to complete
```

```
range_request_timeout_in_ms: 10000
# How long the coordinator should wait for writes to complete
write_request_timeout_in_ms: 2000
# How long the coordinator should wait for counter writes to complete
counter_write_request_timeout_in_ms: 5000
# How long a coordinator should continue to retry a CAS operation
# that contends with other proposals for the same row
cas_contention_timeout_in_ms: 1000
# How long the coordinator should wait for truncates to complete
# (This can be much longer, because unless auto_snapshot is disabled
# we need to flush first so we can snapshot before removing the data.)
truncate_request_timeout_in_ms: 60000
# The default timeout for other, miscellaneous operations
request_timeout_in_ms: 10000

# Enable operation timeout information exchange between nodes to accurately
# measure request timeouts.  If disabled, replicas will assume that requests
# were forwarded to them instantly by the coordinator, which means that
# under overload conditions we will waste that much extra time processing.
# already-timed-out requests.
#
# Warning: before enabling this property make sure to ntp is installed
# and the times are synchronized between the nodes.
cross_node_timeout: false

# Enable socket timeout for streaming operation.
# When a timeout occurs during streaming, streaming is retried from the start
# of the current file. This _can_ involve re-streaming an important amount of
# data, so you should avoid setting the value too low.
# Default value is 3600000, which means streams timeout after an hour.
# streaming_socket_timeout_in_ms: 3600000

# phi value that must be reached for a host to be marked down.
# most users should never need to adjust this.
# phi_convict_threshold: 8

# endpoint_snitch -- Set this to a class that implements
# IEndpointSnitch.  The snitch has two functions:
# - it teaches Cassandra enough about your network topology to route
#   requests efficiently
# - it allows Cassandra to spread replicas around your cluster to avoid
#   correlated failures. It does this by grouping machines into
#   "datacenters" and "racks."  Cassandra will do its best not to have
#   more than one replica on the same "rack" (which may not actually
#   be a physical location)
#
# IF YOU CHANGE THE SNITCH AFTER DATA IS INSERTED INTO THE CLUSTER,
# YOU MUST RUN A FULL REPAIR, SINCE THE SNITCH AFFECTS WHERE REPLICAS
# ARE PLACED.
#
# Out of the box, Cassandra provides
#  - SimpleSnitch:
#    Treats Strategy order as proximity. This can improve cache
#    locality when disabling read repair.  Only appropriate for
#    single-datacenter deployments.
#  - GossipingPropertyFileSnitch
#    This should be your go-to snitch for production use.  The rack
#    and datacenter for the local node are defined in
#    cassandra-rackdc.properties and propagated to other nodes via
#    gossip.  If cassandra-topology.properties exists, it is used as a
#    fallback, allowing migration from the PropertyFileSnitch.
#  - PropertyFileSnitch:
#    Proximity is determined by rack and data center, which are
#    explicitly configured in cassandra-topology.properties.
#  - Ec2Snitch:
#    Appropriate for EC2 deployments in a single Region. Loads Region
#    and Availability Zone information from the EC2 API. The Region is
#    treated as the datacenter, and the Availability Zone as the rack.
#    Only private IPs are used, so this will not work across multiple
#    Regions.
#  - Ec2MultiRegionSnitch:
#    Uses public IPs as broadcast_address to allow cross-region
#    connectivity.  (Thus, you should set seed addresses to the public
```

```
#    IP as well.) You will need to open the storage_port or
#    ssl_storage_port on the public IP firewall.  (For intra-Region
#    traffic, Cassandra will switch to the private IP after
#    establishing a connection.)
#  - RackInferringSnitch:
#    Proximity is determined by rack and data center, which are
#    assumed to correspond to the 3rd and 2nd octet of each node's IP
#    address, respectively.  Unless this happens to match your
#    deployment conventions, this is best used as an example of
#    writing a custom Snitch class and is provided in that spirit.
#
# You can use a custom Snitch by setting this to the full class name
# of the snitch, which will be assumed to be on your classpath.
endpoint_snitch: GossipingPropertyFileSnitch

# controls how often to perform the more expensive part of host score
# calculation
dynamic_snitch_update_interval_in_ms: 100.
# controls how often to reset all host scores, allowing a bad host to
# possibly recover
dynamic_snitch_reset_interval_in_ms: 600000
# if set greater than zero and read_repair_chance is < 1.0, this will allow
# 'pinning' of replicas to hosts in order to increase cache capacity.
# The badness threshold will control how much worse the pinned host has to be
# before the dynamic snitch will prefer other replicas over it.  This is
# expressed as a double which represents a percentage.  Thus, a value of
# 0.2 means Cassandra would continue to prefer the static snitch values
# until the pinned host was 20% worse than the fastest.
dynamic_snitch_badness_threshold: 0.1

# request_scheduler -- Set this to a class that implements
# RequestScheduler, which will schedule incoming client requests
# according to the specific policy. This is useful for multi-tenancy
# with a single Cassandra cluster.
# NOTE: This is specifically for requests from the client and does
# not affect inter node communication.
# org.apache.cassandra.scheduler.NoScheduler - No scheduling takes place
# org.apache.cassandra.scheduler.RoundRobinScheduler - Round robin of
# client requests to a node with a separate queue for each
# request_scheduler_id. The scheduler is further customized by
# request_scheduler_options as described below.
request_scheduler: org.apache.cassandra.scheduler.NoScheduler

# Scheduler Options vary based on the type of scheduler
# NoScheduler - Has no options
# RoundRobin
#  - throttle_limit -- The throttle_limit is the number of in-flight
#                      requests per client.  Requests beyond.
#                      that limit are queued up until
#                      running requests can complete.
#                      The value of 80 here is twice the number of
#                      concurrent_reads + concurrent_writes.
#  - default_weight -- default_weight is optional and allows for
#                      overriding the default which is 1.
#  - weights -- Weights are optional and will default to 1 or the
#             overridden default_weight. The weight translates into how
#             many requests are handled during each turn of the
#             RoundRobin, based on the scheduler id.
#
# request_scheduler_options:
#    throttle_limit: 80
#    default_weight: 5
#    weights:
#      Keyspace1: 1
#      Keyspace2: 5

# request_scheduler_id -- An identifier based on which to perform
# the request scheduling. Currently the only valid option is keyspace.
# request_scheduler_id: keyspace

# Enable or disable inter-node encryption
# Default settings are TLS v1, RSA 1024-bit keys (it is imperative that
# users generate their own keys) TLS_RSA_WITH_AES_128_CBC_SHA as the cipher
```

```
# suite for authentication, key exchange and encryption of the actual data transfers.
# Use the DHE/ECDHE ciphers if running in FIPS 140 compliant mode.
# NOTE: No custom encryption options are enabled at the moment
# The available internode options are : all, none, dc, rack
#
# If set to dc cassandra will encrypt the traffic between the DCs
# If set to rack cassandra will encrypt the traffic between the racks
#
# The passwords used in these options must match the passwords used when generating
# the keystore and truststore.  For instructions on generating these files, see:
# http://download.oracle.com/javase/6/docs/technotes/guides/security/jsse/JSSERefGuide.html#CreateKeystore
#
server_encryption_options:
   internode_encryption: none
   keystore: conf/.keystore
   keystore_password: cassandra
   truststore: conf/.truststore
   truststore_password: cassandra
   # More advanced defaults below:
   # protocol: TLS
   # algorithm: SunX509
   # store_type: JKS
   # cipher_suites:
[TLS_RSA_WITH_AES_128_CBC_SHA,TLS_RSA_WITH_AES_256_CBC_SHA,TLS_DHE_RSA_WITH_AES_
128_CBC_SHA,TLS_DHE_RSA_WITH_AES_256_CBC_SHA,TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,
TLS_ECDHE_RSA_W
   # require_client_auth: false

# enable or disable client/server encryption.
client_encryption_options:
   enabled: false
   keystore: conf/.keystore
   keystore_password: cassandra
   # require_client_auth: false
   # Set trustore and truststore_password if require_client_auth is true
   # truststore: conf/.truststore
   # truststore_password: cassandra
   # More advanced defaults below:
   # protocol: TLS
   # algorithm: SunX509
   # store_type: JKS
   # cipher_suites:
[TLS_RSA_WITH_AES_128_CBC_SHA,TLS_RSA_WITH_AES_256_CBC_SHA,TLS_DHE_RSA_WITH_AES_
128_CBC_SHA,TLS_DHE_RSA_WITH_AES_256_CBC_SHA,TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,
TLS_ECDHE_RSA_W

# internode_compression controls whether traffic between nodes is
# compressed.
# can be:  all  - all traffic is compressed
#          dc   - traffic between different datacenters is compressed
#          none - nothing is compressed.
internode_compression: all

# Enable or disable tcp_nodelay for inter-dc communication.
# Disabling it will result in larger (but fewer) network packets being sent,
# reducing overhead from the TCP protocol itself, at the cost of increasing
# latency if you block for cross-datacenter responses.
inter_dc_tcp_nodelay: false

# GC Pauses greater than gc_warn_threshold_in_ms will be logged at WARN level
# Adjust the threshold based on your application throughput requirement
# By default, Cassandra logs GC Pauses greater than 200 ms at INFO level
# gc_warn_threshold_in_ms: 1000
#
#
#auto_bootstrap: false
```

I. Nazwa serwera: cassandra3

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	cassandra3
IP / Domena	10.90.54.153
OS	CentOS 7
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	100 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Bazodanowy

2. Konfiguracja .hosts

```
127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6
```

3. Interfejsy sieciowe

eth0 10.90.54.153

4. Zainstalowane usługi

a. Cassandra

Parametr	Wartość
Lokalizacja instalacji	/etc/cassandra
Zawartość pliku konfiguracyjnego #1	/etc/cassandra/conf # Cassandra storage config YAML. # NOTE: # See http://wiki.apache.org/cassandra/StorageConfiguration for # full explanations of configuration directives # /NOTE # The name of the cluster. This is mainly used to prevent machines in # one logical cluster from joining another. cluster_name: 'EpoCluster' # This defines the number of tokens randomly assigned to this node on the ring # The more tokens, relative to other nodes, the larger the proportion of data # that this node will store. You probably want all nodes to have the same number # of tokens assuming they have equal hardware capability. # # If you leave this unspecified, Cassandra will use the default of 1 token for legacy compatibility, # and will use the initial_token as described below. # # Specifying initial_token will override this setting on the node's initial start, # on subsequent starts, this setting will apply even if initial token is set. #

```
# If you already have a cluster with 1 token per node, and wish to migrate to.
# multiple tokens per node, see http://wiki.apache.org/cassandra/Operations
num_tokens: 256

# initial_token allows you to specify tokens manually.  While you can use # it with
# vnodes (num_tokens > 1, above) -- in which case you should provide a.
# comma-separated list -- it's primarily used when adding nodes # to legacy clusters.
# that do not have vnodes enabled.
# initial_token:

# See http://wiki.apache.org/cassandra/HintedHandoff
# May either be "true" or "false" to enable globally, or contain a list
# of data centers to enable per-datacenter.
# hinted_handoff_enabled: DC1,DC2
hinted_handoff_enabled: true
# this defines the maximum amount of time a dead host will have hints
# generated.  After it has been dead this long, new hints for it will not be
# created until it has been seen alive and gone down again.
max_hint_window_in_ms: 10800000 # 3 hours
# Maximum throttle in KBs per second, per delivery thread.  This will be
# reduced proportionally to the number of nodes in the cluster.  (If there
# are two nodes in the cluster, each delivery thread will use the maximum
# rate; if there are three, each will throttle to half of the maximum,
# since we expect two nodes to be delivering hints simultaneously.)
hinted_handoff_throttle_in_kb: 1024
# Number of threads with which to deliver hints;
# Consider increasing this number when you have multi-dc deployments, since
# cross-dc handoff tends to be slower
max_hints_delivery_threads: 2

# Maximum throttle in KBs per second, total. This will be
# reduced proportionally to the number of nodes in the cluster.
batchlog_replay_throttle_in_kb: 1024

# Authentication backend, implementing IAuthenticator; used to identify users
# Out of the box, Cassandra provides org.apache.cassandra.auth.{AllowAllAuthenticator,
# PasswordAuthenticator}.
#
# - AllowAllAuthenticator performs no checks - set it to disable authentication.
# - PasswordAuthenticator relies on username/password pairs to authenticate
#   users. It keeps usernames and hashed passwords in system_auth.credentials table.
#   Please increase system_auth keyspace replication factor if you use this authenticator.
authenticator: PasswordAuthenticator

# Authorization backend, implementing IAuthorizer; used to limit access/provide permissions
# Out of the box, Cassandra provides org.apache.cassandra.auth.{AllowAllAuthorizer,
# CassandraAuthorizer}.
#
# - AllowAllAuthorizer allows any action to any user - set it to disable authorization.
# - CassandraAuthorizer stores permissions in system_auth.permissions table. Please
#   increase system_auth keyspace replication factor if you use this authorizer.
authorizer: CassandraAuthorizer

# Validity period for permissions cache (fetching permissions can be an
# expensive operation depending on the authorizer, CassandraAuthorizer is
# one example). Defaults to 2000, set to 0 to disable.
# Will be disabled automatically for AllowAllAuthorizer.
permissions_validity_in_ms: 2000

# Refresh interval for permissions cache (if enabled).
# After this interval, cache entries become eligible for refresh. Upon next
# access, an async reload is scheduled and the old value returned until it
# completes. If permissions_validity_in_ms is non-zero, then this must be
# also.
# Defaults to the same value as permissions_validity_in_ms.
# permissions_update_interval_in_ms: 1000

# The partitioner is responsible for distributing groups of rows (by
# partition key) across nodes in the cluster.  You should leave this
# alone for new clusters.  The partitioner can NOT be changed without
# reloading all data, so when upgrading you should set this to the
# same partitioner you were already using.
#
```

```
# Besides Murmur3Partitioner, partitioners included for backwards
# compatibility include RandomPartitioner, ByteOrderedPartitioner, and
# OrderPreservingPartitioner.
#
partitioner: org.apache.cassandra.dht.Murmur3Partitioner

# Directories where Cassandra should store data on disk.  Cassandra
# will spread data evenly across them, subject to the granularity of
# the configured compaction strategy.
# If not set, the default directory is $CASSANDRA_HOME/data/data.
data_file_directories:
    - /opt/cassandra/data

# commit log.  when running on magnetic HDD, this should be a
# separate spindle than the data directories.
# If not set, the default directory is $CASSANDRA_HOME/data/commitlog.
commitlog_directory: /opt/cassandra/commitlog

# policy for data disk failures:
# die: shut down gossip and client transports and kill the JVM for any fs errors or
#      single-sstable errors, so the node can be replaced.
# stop_paranoid: shut down gossip and client transports even for single-sstable errors,
#               kill the JVM for errors during startup.
# stop: shut down gossip and client transports, leaving the node effectively dead, but
#       can still be inspected via JMX, kill the JVM for errors during startup.
# best_effort: stop using the failed disk and respond to requests based on
#              remaining available sstables.  This means you WILL see obsolete
#              data at CL.ONE!
# ignore: ignore fatal errors and let requests fail, as in pre-1.2 Cassandra
disk_failure_policy: stop

# policy for commit disk failures:
# die: shut down gossip and Thrift and kill the JVM, so the node can be replaced.
# stop: shut down gossip and Thrift, leaving the node effectively dead, but
#       can still be inspected via JMX.
# stop_commit: shutdown the commit log, letting writes collect but
#              continuing to service reads, as in pre-2.0.5 Cassandra
# ignore: ignore fatal errors and let the batches fail
commit_failure_policy: stop

# Maximum size of the key cache in memory.
#
# Each key cache hit saves 1 seek and each row cache hit saves 2 seeks at the
# minimum, sometimes more. The key cache is fairly tiny for the amount of
# time it saves, so it's worthwhile to use it at large numbers.
# The row cache saves even more time, but must contain the entire row,
# so it is extremely space-intensive. It's best to only use the
# row cache if you have hot rows or static rows.
#
# NOTE: if you reduce the size, you may not get you hottest keys loaded on startup.
#
# Default value is empty to make it "auto" (min(5% of Heap (in MB), 100MB)). Set to 0 to disable key cache.
key_cache_size_in_mb:

# Duration in seconds after which Cassandra should
# save the key cache. Caches are saved to saved_caches_directory as
# specified in this configuration file.
#
# Saved caches greatly improve cold-start speeds, and is relatively cheap in
# terms of I/O for the key cache. Row cache saving is much more expensive and
# has limited use.
#
# Default is 14400 or 4 hours.
key_cache_save_period: 14400

# Number of keys from the key cache to save
# Disabled by default, meaning all keys are going to be saved
# key_cache_keys_to_save: 100

# Maximum size of the row cache in memory.
# NOTE: if you reduce the size, you may not get you hottest keys loaded on startup.
#
# Default value is 0, to disable row caching.
```

```
row_cache_size_in_mb: 0

# Duration in seconds after which Cassandra should
# save the row cache. Caches are saved to saved_caches_directory as specified
# in this configuration file.
#
# Saved caches greatly improve cold-start speeds, and is relatively cheap in
# terms of I/O for the key cache. Row cache saving is much more expensive and
# has limited use.
#
# Default is 0 to disable saving the row cache.
row_cache_save_period: 0

# Number of keys from the row cache to save
# Disabled by default, meaning all keys are going to be saved
# row_cache_keys_to_save: 100

# Maximum size of the counter cache in memory.
#
# Counter cache helps to reduce counter locks' contention for hot counter cells.
# In case of RF = 1 a counter cache hit will cause Cassandra to skip the read before
# write entirely. With RF > 1 a counter cache hit will still help to reduce the duration
# of the lock hold, helping with hot counter cell updates, but will not allow skipping
# the read entirely. Only the local (clock, count) tuple of a counter cell is kept
# in memory, not the whole counter, so it's relatively cheap.
#
# NOTE: if you reduce the size, you may not get you hottest keys loaded on startup.
#
# Default value is empty to make it "auto" (min(2.5% of Heap (in MB), 50MB)). Set to 0 to disable counter cache.
# NOTE: if you perform counter deletes and rely on low gcgs, you should disable the counter cache.
counter_cache_size_in_mb:

# Duration in seconds after which Cassandra should
# save the counter cache (keys only). Caches are saved to saved_caches_directory as
# specified in this configuration file.
#
# Default is 7200 or 2 hours.
counter_cache_save_period: 7200

# Number of keys from the counter cache to save
# Disabled by default, meaning all keys are going to be saved
# counter_cache_keys_to_save: 100

# The off-heap memory allocator.  Affects storage engine metadata as
# well as caches.  Experiments show that JEMAlloc saves some memory
# than the native GCC allocator (i.e., JEMalloc is more
# fragmentation-resistant).
#.
# Supported values are: NativeAllocator, JEMallocAllocator
#
# If you intend to use JEMallocAllocator you have to install JEMalloc as library and
# modify cassandra-env.sh as directed in the file.
#
# Defaults to NativeAllocator
# memory_allocator: NativeAllocator

# saved caches
# If not set, the default directory is $CASSANDRA_HOME/data/saved_caches.
saved_caches_directory: /opt/cassandra/saved_caches

# commitlog_sync may be either "periodic" or "batch.".
#.
# When in batch mode, Cassandra won't ack writes until the commit log
# has been fsynced to disk.  It will wait
# commitlog_sync_batch_window_in_ms milliseconds between fsyncs.
# This window should be kept short because the writer threads will
# be unable to do extra work while waiting.  (You may need to increase
# concurrent_writes for the same reason.)
#
# commitlog_sync: batch
# commitlog_sync_batch_window_in_ms: 2
#
# the other option is "periodic" where writes may be acked immediately
```

```
# and the CommitLog is simply synced every commitlog_sync_period_in_ms
# milliseconds..
commitlog_sync: periodic
commitlog_sync_period_in_ms: 10000

# The size of the individual commitlog file segments.  A commitlog
# segment may be archived, deleted, or recycled once all the data
# in it (potentially from each columnfamily in the system) has been
# flushed to sstables...
#
# The default size is 32, which is almost always fine, but if you are
# archiving commitlog segments (see commitlog_archiving.properties),
# then you probably want a finer granularity of archiving; 8 or 16 MB
# is reasonable.
commitlog_segment_size_in_mb: 32

# Reuse commit log files when possible. The default is false, and this
# feature will be removed entirely in future versions of Cassandra.
#commitlog_segment_recycling: false

# any class that implements the SeedProvider interface and has a
# constructor that takes a Map<String, String> of parameters will do.
seed_provider:
   # Addresses of hosts that are deemed contact points..
   # Cassandra nodes use this list of hosts to find each other and learn
   # the topology of the ring.  You must change this if you are running
   # multiple nodes!
   - class_name: org.apache.cassandra.locator.SimpleSeedProvider
     parameters:
         # seeds is actually a comma-delimited list of addresses.
         # Ex: "<ip1>,<ip2>,<ip3>"
         - seeds: "10.90.54.151,10.90.54.152,10.90.54.153,10.90.54.154"

# For workloads with more data than can fit in memory, Cassandra's
# bottleneck will be reads that need to fetch data from
# disk. "concurrent_reads" should be set to (16 * number_of_drives) in
# order to allow the operations to enqueue low enough in the stack
# that the OS and drives can reorder them. Same applies to
# "concurrent_counter_writes", since counter writes read the current
# values before incrementing and writing them back.
#
# On the other hand, since writes are almost never IO bound, the ideal
# number of "concurrent_writes" is dependent on the number of cores in
# your system; (8 * number_of_cores) is a good rule of thumb.
concurrent_reads: 32
concurrent_writes: 32
concurrent_counter_writes: 32

# Total memory to use for sstable-reading buffers.  Defaults to
# the smaller of 1/4 of heap or 512MB.
# file_cache_size_in_mb: 512

# Total permitted memory to use for memtables. Cassandra will stop.
# accepting writes when the limit is exceeded until a flush completes,
# and will trigger a flush based on memtable_cleanup_threshold
# If omitted, Cassandra will set both to 1/4 the size of the heap.
# memtable_heap_space_in_mb: 2048
# memtable_offheap_space_in_mb: 2048

# Ratio of occupied non-flushing memtable size to total permitted size
# that will trigger a flush of the largest memtable.  Lager mct will
# mean larger flushes and hence less compaction, but also less concurrent
# flush activity which can make it difficult to keep your disks fed
# under heavy write load.
#
# memtable_cleanup_threshold defaults to 1 / (memtable_flush_writers + 1)
# memtable_cleanup_threshold: 0.11

# Specify the way Cassandra allocates and manages memtable memory.
# Options are:
#   heap_buffers:    on heap nio buffers
#   offheap_buffers: off heap (direct) nio buffers
#   offheap_objects: native memory, eliminating nio buffer heap overhead
```

```
memtable_allocation_type: heap_buffers

# Total space to use for commitlogs.  Since commitlog segments are
# mmapped, and hence use up address space, the default size is 32
# on 32-bit JVMs, and 8192 on 64-bit JVMs.
#
# If space gets above this value (it will round up to the next nearest
# segment multiple), Cassandra will flush every dirty CF in the oldest
# segment and remove it.  So a small total commitlog space will tend
# to cause more flush activity on less-active columnfamilies.
# commitlog_total_space_in_mb: 8192

# This sets the amount of memtable flush writer threads.  These will
# be blocked by disk io, and each one will hold a memtable in memory
# while blocked..
#
# memtable_flush_writers defaults to the smaller of (number of disks,
# number of cores), with a minimum of 2 and a maximum of 8.
#.
# If your data directories are backed by SSD, you should increase this
# to the number of cores.
#memtable_flush_writers: 8

# A fixed memory pool size in MB for for SSTable index summaries. If left
# empty, this will default to 5% of the heap size. If the memory usage of
# all index summaries exceeds this limit, SSTables with low read rates will
# shrink their index summaries in order to meet this limit.  However, this
# is a best-effort process. In extreme conditions Cassandra may need to use
# more than this amount of memory.
index_summary_capacity_in_mb:

# How frequently index summaries should be resampled.  This is done
# periodically to redistribute memory from the fixed-size pool to sstables
# proportional their recent read rates.  Setting to -1 will disable this
# process, leaving existing index summaries at their current sampling level.
index_summary_resize_interval_in_minutes: 60

# Whether to, when doing sequential writing, fsync() at intervals in
# order to force the operating system to flush the dirty
# buffers. Enable this to avoid sudden dirty buffer flushing from
# impacting read latencies. Almost always a good idea on SSDs; not
# necessarily on platters.
trickle_fsync: false
trickle_fsync_interval_in_kb: 10240

# TCP port, for commands and data
# For security reasons, you should not expose this port to the internet.  Firewall it if needed.
storage_port: 7000

# SSL port, for encrypted communication.  Unused unless enabled in
# encryption_options
# For security reasons, you should not expose this port to the internet.  Firewall it if needed.
ssl_storage_port: 7001

# Address or interface to bind to and tell other Cassandra nodes to connect to.
# You _must_ change this if you want multiple nodes to be able to communicate!
#
# Set listen_address OR listen_interface, not both. Interfaces must correspond
# to a single address, IP aliasing is not supported.
#
# Leaving it blank leaves it up to InetAddress.getLocalHost(). This
# will always do the Right Thing _if_ the node is properly configured
# (hostname, name resolution, etc), and the Right Thing is to use the
# address associated with the hostname (it might not be).
#
# Setting listen_address to 0.0.0.0 is always wrong.
#
# If you choose to specify the interface by name and the interface has an ipv4 and an ipv6 address
# you can specify which should be chosen using listen_interface_prefer_ipv6. If false the first ipv4
# address will be used. If true the first ipv6 address will be used. Defaults to false preferring
# ipv4. If there is only one address it will be selected regardless of ipv4/ipv6.
listen_address: 10.90.54.151
# listen_interface: eth0
```

```
# listen_interface_prefer_ipv6: false

# Address to broadcast to other Cassandra nodes
# Leaving this blank will set it to the same value as listen_address
# broadcast_address: 1.2.3.4

# Internode authentication backend, implementing IInternodeAuthenticator;
# used to allow/disallow connections from peer nodes.
# internode_authenticator: org.apache.cassandra.auth.AllowAllInternodeAuthenticator

# Whether to start the native transport server.
# Please note that the address on which the native transport is bound is the
# same as the rpc_address. The port however is different and specified below.
start_native_transport: true
# port for the CQL native transport to listen for clients on
# For security reasons, you should not expose this port to the internet.  Firewall it if needed.
native_transport_port: 9042
# The maximum threads for handling requests when the native transport is used.
# This is similar to rpc_max_threads though the default differs slightly (and
# there is no native_transport_min_threads, idle threads will always be stopped
# after 30 seconds).
# native_transport_max_threads: 128
#
# The maximum size of allowed frame. Frame (requests) larger than this will
# be rejected as invalid. The default is 256MB.
# native_transport_max_frame_size_in_mb: 256

# The maximum number of concurrent client connections.
# The default is -1, which means unlimited.
# native_transport_max_concurrent_connections: -1

# The maximum number of concurrent client connections per source ip.
# The default is -1, which means unlimited.
# native_transport_max_concurrent_connections_per_ip: -1


# Whether to start the thrift rpc server.
start_rpc: true

# The address or interface to bind the Thrift RPC service and native transport
# server to.
#
# Set rpc_address OR rpc_interface, not both. Interfaces must correspond
# to a single address, IP aliasing is not supported.
#
# Leaving rpc_address blank has the same effect as on listen_address
# (i.e. it will be based on the configured hostname of the node).
#
# Note that unlike listen_address, you can specify 0.0.0.0, but you must also
# set broadcast_rpc_address to a value other than 0.0.0.0.
#
# For security reasons, you should not expose this port to the internet.  Firewall it if needed.
#
# If you choose to specify the interface by name and the interface has an ipv4 and an ipv6 address
# you can specify which should be chosen using rpc_interface_prefer_ipv6. If false the first ipv4
# address will be used. If true the first ipv6 address will be used. Defaults to false preferring
# ipv4. If there is only one address it will be selected regardless of ipv4/ipv6.
rpc_address: 10.90.54.151
# rpc_interface: eth1
# rpc_interface_prefer_ipv6: false

# port for Thrift to listen for clients on
rpc_port: 9160

# RPC address to broadcast to drivers and other Cassandra nodes. This cannot
# be set to 0.0.0.0. If left blank, this will be set to the value of
# rpc_address. If rpc_address is set to 0.0.0.0, broadcast_rpc_address must
# be set.
# broadcast_rpc_address: 1.2.3.4

# enable or disable keepalive on rpc/native connections
rpc_keepalive: true
```

```
# Cassandra provides two out-of-the-box options for the RPC Server:
#
# sync  -> One thread per thrift connection. For a very large number of clients, memory
#          will be your limiting factor. On a 64 bit JVM, 180KB is the minimum stack size
#          per thread, and that will correspond to your use of virtual memory (but physical memory
#          may be limited depending on use of stack space).
#
# hsha  -> Stands for "half synchronous, half asynchronous." All thrift clients are handled
#          asynchronously using a small number of threads that does not vary with the amount
#          of thrift clients (and thus scales well to many clients). The rpc requests are still
#          synchronous (one thread per active request). If hsha is selected then it is essential
#          that rpc_max_threads is changed from the default value of unlimited.
#
# The default is sync because on Windows hsha is about 30% slower.  On Linux,
# sync/hsha performance is about the same, with hsha of course using less memory.
#
# Alternatively,  can provide your own RPC server by providing the fully-qualified class name
# of an o.a.c.t.TServerFactory that can create an instance of it.
rpc_server_type: sync

# Uncomment rpc_min|max_thread to set request pool size limits.
#
# Regardless of your choice of RPC server (see above), the number of maximum requests in the
# RPC thread pool dictates how many concurrent requests are possible (but if you are using the sync
# RPC server, it also dictates the number of clients that can be connected at all).
#
# The default is unlimited and thus provides no protection against clients overwhelming the server. You are
# encouraged to set a maximum that makes sense for you in production, but do keep in mind that
# rpc_max_threads represents the maximum number of client requests this server may execute concurrently.
#
# rpc_min_threads: 16
# rpc_max_threads: 2048

# uncomment to set socket buffer sizes on rpc connections
# rpc_send_buff_size_in_bytes:
# rpc_recv_buff_size_in_bytes:

# Uncomment to set socket buffer size for internode communication
# Note that when setting this, the buffer size is limited by net.core.wmem_max
# and when not setting it it is defined by net.ipv4.tcp_wmem
# See:
# /proc/sys/net/core/wmem_max
# /proc/sys/net/core/rmem_max
# /proc/sys/net/ipv4/tcp_wmem
# /proc/sys/net/ipv4/tcp_wmem
# and: man tcp
# internode_send_buff_size_in_bytes:
# internode_recv_buff_size_in_bytes:

# Frame size for thrift (maximum message length).
thrift_framed_transport_size_in_mb: 15

# Set to true to have Cassandra create a hard link to each sstable
# flushed or streamed locally in a backups/ subdirectory of the
# keyspace data.  Removing these links is the operator's
# responsibility.
incremental_backups: false

# Whether or not to take a snapshot before each compaction.  Be
# careful using this option, since Cassandra won't clean up the
# snapshots for you.  Mostly useful if you're paranoid when there
# is a data format change.
snapshot_before_compaction: false

# Whether or not a snapshot is taken of the data before keyspace truncation
# or dropping of column families. The STRONGLY advised default of true.
# should be used to provide data safety. If you set this flag to false, you will
# lose data on truncation or drop.
auto_snapshot: true

# When executing a scan, within or across a partition, we need to keep the
# tombstones seen in memory so we can return them to the coordinator, which
# will use them to make sure other replicas also know about the deleted rows.
```

```
# With workloads that generate a lot of tombstones, this can cause performance
# problems and even exaust the server heap.
# (http://www.datastax.com/dev/blog/cassandra-anti-patterns-queues-and-queue-like-datasets)
# Adjust the thresholds here if you understand the dangers and want to
# scan more tombstones anyway.  These thresholds may also be adjusted at runtime
# using the StorageService mbean.
tombstone_warn_threshold: 1000
tombstone_failure_threshold: 100000

# Granularity of the collation index of rows within a partition.
# Increase if your rows are large, or if you have a very large
# number of rows per partition.  The competing goals are these:
#   1) a smaller granularity means more index entries are generated
#      and looking up rows withing the partition by collation column
#      is faster
#   2) but, Cassandra will keep the collation index in memory for hot
#      rows (as part of the key cache), so a larger granularity means
#      you can cache more hot rows
column_index_size_in_kb: 64


# Log WARN on any batch size exceeding this value. 5kb per batch by default.
# Caution should be taken on increasing the size of this threshold as it can lead to node instability.
batch_size_warn_threshold_in_kb: 5

# Number of simultaneous compactions to allow, NOT including
# validation "compactions" for anti-entropy repair.  Simultaneous
# compactions can help preserve read performance in a mixed read/write
# workload, by mitigating the tendency of small sstables to accumulate
# during a single long running compactions. The default is usually
# fine and if you experience problems with compaction running too
# slowly or too fast, you should look at
# compaction_throughput_mb_per_sec first.
#
# concurrent_compactors defaults to the smaller of (number of disks,
# number of cores), with a minimum of 2 and a maximum of 8.
#.
# If your data directories are backed by SSD, you should increase this
# to the number of cores.
#concurrent_compactors: 1

# Throttles compaction to the given total throughput across the entire
# system. The faster you insert data, the faster you need to compact in
# order to keep the sstable count down, but in general, setting this to
# 16 to 32 times the rate you are inserting data is more than sufficient.
# Setting this to 0 disables throttling. Note that this account for all types
# of compaction, including validation compaction.
compaction_throughput_mb_per_sec: 16

# Log a warning when compacting partitions larger than this value
compaction_large_partition_warning_threshold_mb: 100

# When compacting, the replacement sstable(s) can be opened before they
# are completely written, and used in place of the prior sstables for
# any range that has been written. This helps to smoothly transfer reads.
# between the sstables, reducing page cache churn and keeping hot rows hot
sstable_preemptive_open_interval_in_mb: 50

# Throttles all outbound streaming file transfers on this node to the
# given total throughput in Mbps. This is necessary because Cassandra does
# mostly sequential IO when streaming data during bootstrap or repair, which
# can lead to saturating the network connection and degrading rpc performance.
# When unset, the default is 200 Mbps or 25 MB/s.
# stream_throughput_outbound_megabits_per_sec: 200

# Throttles all streaming file transfer between the datacenters,
# this setting allows users to throttle inter dc stream throughput in addition
# to throttling all network stream traffic as configured with
# stream_throughput_outbound_megabits_per_sec
# inter_dc_stream_throughput_outbound_megabits_per_sec:

# How long the coordinator should wait for read operations to complete
read_request_timeout_in_ms: 5000
```

```
# How long the coordinator should wait for seq or index scans to complete
range_request_timeout_in_ms: 10000
# How long the coordinator should wait for writes to complete
write_request_timeout_in_ms: 2000
# How long the coordinator should wait for counter writes to complete
counter_write_request_timeout_in_ms: 5000
# How long a coordinator should continue to retry a CAS operation
# that contends with other proposals for the same row
cas_contention_timeout_in_ms: 1000
# How long the coordinator should wait for truncates to complete
# (This can be much longer, because unless auto_snapshot is disabled
# we need to flush first so we can snapshot before removing the data.)
truncate_request_timeout_in_ms: 60000
# The default timeout for other, miscellaneous operations
request_timeout_in_ms: 10000

# Enable operation timeout information exchange between nodes to accurately
# measure request timeouts.  If disabled, replicas will assume that requests
# were forwarded to them instantly by the coordinator, which means that
# under overload conditions we will waste that much extra time processing.
# already-timed-out requests.
#
# Warning: before enabling this property make sure to ntp is installed
# and the times are synchronized between the nodes.
cross_node_timeout: false

# Enable socket timeout for streaming operation.
# When a timeout occurs during streaming, streaming is retried from the start
# of the current file. This _can_ involve re-streaming an important amount of
# data, so you should avoid setting the value too low.
# Default value is 3600000, which means streams timeout after an hour.
# streaming_socket_timeout_in_ms: 3600000

# phi value that must be reached for a host to be marked down.
# most users should never need to adjust this.
# phi_convict_threshold: 8

# endpoint_snitch -- Set this to a class that implements
# IEndpointSnitch.  The snitch has two functions:
# - it teaches Cassandra enough about your network topology to route
#   requests efficiently
# - it allows Cassandra to spread replicas around your cluster to avoid
#   correlated failures. It does this by grouping machines into
#   "datacenters" and "racks."  Cassandra will do its best not to have
#   more than one replica on the same "rack" (which may not actually
#   be a physical location)
#
# IF YOU CHANGE THE SNITCH AFTER DATA IS INSERTED INTO THE CLUSTER,
# YOU MUST RUN A FULL REPAIR, SINCE THE SNITCH AFFECTS WHERE REPLICAS
# ARE PLACED.
#
# Out of the box, Cassandra provides
#  - SimpleSnitch:
#    Treats Strategy order as proximity. This can improve cache
#    locality when disabling read repair.  Only appropriate for
#    single-datacenter deployments.
#  - GossipingPropertyFileSnitch
#    This should be your go-to snitch for production use.  The rack
#    and datacenter for the local node are defined in
#    cassandra-rackdc.properties and propagated to other nodes via
#    gossip.  If cassandra-topology.properties exists, it is used as a
#    fallback, allowing migration from the PropertyFileSnitch.
#  - PropertyFileSnitch:
#    Proximity is determined by rack and data center, which are
#    explicitly configured in cassandra-topology.properties.
#  - Ec2Snitch:
#    Appropriate for EC2 deployments in a single Region. Loads Region
#    and Availability Zone information from the EC2 API. The Region is
#    treated as the datacenter, and the Availability Zone as the rack.
#    Only private IPs are used, so this will not work across multiple
#    Regions.
#  - Ec2MultiRegionSnitch:
#    Uses public IPs as broadcast_address to allow cross-region
```

```
#    connectivity.  (Thus, you should set seed addresses to the public
#    IP as well.) You will need to open the storage_port or
#    ssl_storage_port on the public IP firewall.  (For intra-Region
#    traffic, Cassandra will switch to the private IP after
#    establishing a connection.)
#  - RackInferringSnitch:
#    Proximity is determined by rack and data center, which are
#    assumed to correspond to the 3rd and 2nd octet of each node's IP
#    address, respectively.  Unless this happens to match your
#    deployment conventions, this is best used as an example of
#    writing a custom Snitch class and is provided in that spirit.
#
# You can use a custom Snitch by setting this to the full class name
# of the snitch, which will be assumed to be on your classpath.
endpoint_snitch: GossipingPropertyFileSnitch

# controls how often to perform the more expensive part of host score
# calculation
dynamic_snitch_update_interval_in_ms: 100.
# controls how often to reset all host scores, allowing a bad host to
# possibly recover
dynamic_snitch_reset_interval_in_ms: 600000
# if set greater than zero and read_repair_chance is < 1.0, this will allow
# 'pinning' of replicas to hosts in order to increase cache capacity.
# The badness threshold will control how much worse the pinned host has to be
# before the dynamic snitch will prefer other replicas over it.  This is
# expressed as a double which represents a percentage.  Thus, a value of
# 0.2 means Cassandra would continue to prefer the static snitch values
# until the pinned host was 20% worse than the fastest.
dynamic_snitch_badness_threshold: 0.1

# request_scheduler -- Set this to a class that implements
# RequestScheduler, which will schedule incoming client requests
# according to the specific policy. This is useful for multi-tenancy
# with a single Cassandra cluster.
# NOTE: This is specifically for requests from the client and does
# not affect inter node communication.
# org.apache.cassandra.scheduler.NoScheduler - No scheduling takes place
# org.apache.cassandra.scheduler.RoundRobinScheduler - Round robin of
# client requests to a node with a separate queue for each
# request_scheduler_id. The scheduler is further customized by
# request_scheduler_options as described below.
request_scheduler: org.apache.cassandra.scheduler.NoScheduler

# Scheduler Options vary based on the type of scheduler
# NoScheduler - Has no options
# RoundRobin
#  - throttle_limit -- The throttle_limit is the number of in-flight
#                      requests per client.  Requests beyond.
#                      that limit are queued up until
#                      running requests can complete.
#                      The value of 80 here is twice the number of
#                      concurrent_reads + concurrent_writes.
#  - default_weight -- default_weight is optional and allows for
#                      overriding the default which is 1.
#  - weights -- Weights are optional and will default to 1 or the
#               overridden default_weight. The weight translates into how
#               many requests are handled during each turn of the
#               RoundRobin, based on the scheduler id.
#
# request_scheduler_options:
#    throttle_limit: 80
#    default_weight: 5
#    weights:
#      Keyspace1: 1
#      Keyspace2: 5

# request_scheduler_id -- An identifier based on which to perform
# the request scheduling. Currently the only valid option is keyspace.
# request_scheduler_id: keyspace

# Enable or disable inter-node encryption
# Default settings are TLS v1, RSA 1024-bit keys (it is imperative that
```

```
# users generate their own keys) TLS_RSA_WITH_AES_128_CBC_SHA as the cipher
# suite for authentication, key exchange and encryption of the actual data transfers.
# Use the DHE/ECDHE ciphers if running in FIPS 140 compliant mode.
# NOTE: No custom encryption options are enabled at the moment
# The available internode options are : all, none, dc, rack
#
# If set to dc cassandra will encrypt the traffic between the DCs
# If set to rack cassandra will encrypt the traffic between the racks
#
# The passwords used in these options must match the passwords used when generating
# the keystore and truststore.  For instructions on generating these files, see:
# http://download.oracle.com/javase/6/docs/technotes/guides/security/jsse/JSSERefGuide.html#CreateKeystore
#
server_encryption_options:
    internode_encryption: none
    keystore: conf/.keystore
    keystore_password: cassandra
    truststore: conf/.truststore
    truststore_password: cassandra
    # More advanced defaults below:
    # protocol: TLS
    # algorithm: SunX509
    # store_type: JKS
    # cipher_suites:
[TLS_RSA_WITH_AES_128_CBC_SHA,TLS_RSA_WITH_AES_256_CBC_SHA,TLS_DHE_RSA_WITH_AES_
128_CBC_SHA,TLS_DHE_RSA_WITH_AES_256_CBC_SHA,TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,
TLS_ECDHE_RSA_W
    # require_client_auth: false

# enable or disable client/server encryption.
client_encryption_options:
    enabled: false
    keystore: conf/.keystore
    keystore_password: cassandra
    # require_client_auth: false
    # Set trustore and truststore_password if require_client_auth is true
    # truststore: conf/.truststore
    # truststore_password: cassandra
    # More advanced defaults below:
    # protocol: TLS
    # algorithm: SunX509
    # store_type: JKS
    # cipher_suites:
[TLS_RSA_WITH_AES_128_CBC_SHA,TLS_RSA_WITH_AES_256_CBC_SHA,TLS_DHE_RSA_WITH_AES_
128_CBC_SHA,TLS_DHE_RSA_WITH_AES_256_CBC_SHA,TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,
TLS_ECDHE_RSA_W

# internode_compression controls whether traffic between nodes is
# compressed.
# can be:  all  - all traffic is compressed
#          dc   - traffic between different datacenters is compressed
#          none - nothing is compressed.
internode_compression: all

# Enable or disable tcp_nodelay for inter-dc communication.
# Disabling it will result in larger (but fewer) network packets being sent,
# reducing overhead from the TCP protocol itself, at the cost of increasing
# latency if you block for cross-datacenter responses.
inter_dc_tcp_nodelay: false

# GC Pauses greater than gc_warn_threshold_in_ms will be logged at WARN level
# Adjust the threshold based on your application throughput requirement
# By default, Cassandra logs GC Pauses greater than 200 ms at INFO level
# gc_warn_threshold_in_ms: 1000
#
#
#auto_bootstrap: false
```

I. Nazwa serwera: db1

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	db1
IP / Domena	10.90.54.221
OS	Ubuntu 16.04
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	50
Zewnętrzne IP:	NIE
Funkcja serwera:	bazodanowy

2. Konfiguracja .hosts

```
10.90.54.201 app1
10.90.54.202 app2
10.90.54.211 storage1
10.90.54.212 storage2
10.90.54.221 db1 portal-db1

127.0.0.1 localhost portal-db1


# The following lines are desirable for IPv6 capable hosts
::1 ip6-localhost ip6-loopback
fe00::0 ip6-localnet
ff00::0 ip6-mcastprefix
ff02::1 ip6-allnodes
ff02::2 ip6-allrouters
ff02::3 ip6-allhosts
```

3. Interfejsy sieciowe

ens3 10.90.54.221

4. Zainstalowane usługi

a. Postrgresql

Parametr	Wartość
Lokalizacja instalacji	/etc/postgresql
Zawartość pliku konfiguracyjnego #1	/etc/postgresql/9.5/main/postgresql.conf # ----------------------------- # PostgreSQL configuration file # ----------------------------- # # This file consists of lines of the form: # # name = value # # (The "=" is optional.) Whitespace may be used. Comments are introduced with # "#" anywhere on a line. The complete list of parameter names and allowed # values can be found in the PostgreSQL documentation.

```
#
# The commented-out settings shown in this file represent the default values.
# Re-commenting a setting is NOT sufficient to revert it to the default value;
# you need to reload the server.
#
# This file is read on server startup and when the server receives a SIGHUP
# signal.  If you edit the file on a running system, you have to SIGHUP the
# server for the changes to take effect, or use "pg_ctl reload".  Some
# parameters, which are marked below, require a server shutdown and restart to
# take effect.
#
# Any parameter can also be given as a command-line option to the server, e.g.,
# "postgres -c log_connections=on".  Some parameters can be changed at run time
# with the "SET" SQL command.
#
# Memory units:  kB = kilobytes        Time units:  ms  = milliseconds
#                MB = megabytes                 s   = seconds
#                GB = gigabytes                 min = minutes
#                TB = terabytes                 h   = hours
#                                               d   = days


#------------------------------------------------------------------------------
# FILE LOCATIONS
#------------------------------------------------------------------------------

# The default values of these variables are driven from the -D command-line
# option or PGDATA environment variable, represented here as ConfigDir.

data_directory = '/var/lib/postgresql/9.5/main'><------># use data in another directory
<------><------><------><------><------># (change requires restart)
hba_file = '/etc/postgresql/9.5/main/pg_hba.conf'<----->#host-based authentication file
<------><------><------><------><------># (change requires restart)
ident_file = '/etc/postgresql/9.5/main/pg_ident.conf'<->#ident configuration file
<------><------><------><------><------># (change requires restart)

# If external_pid_file is not explicitly set, no extra PID file is written.
external_pid_file = '/var/run/postgresql/9.5-main.pid'<><------><------># write an extra PID file
<------><------><------><------><------># (change requires restart)


#------------------------------------------------------------------------------
# CONNECTIONS AND AUTHENTICATION
#------------------------------------------------------------------------------

# - Connection Settings -

listen_addresses = 'localhost,db1'<----><------># what IP address(es) to listen on;
<------><------><------><------><------># comma-separated list of addresses;
<------><------><------><------><------># defaults to 'localhost'; use '*' for all
port = 5432<---><------><------><------># (change requires restart)
max_connections = 1000<><------><------># (change requires restart)
#superuser_reserved_connections = 3<---># (change requires restart)
unix_socket_directories = '/var/run/postgresql'># comma-separated list of directories
<------><------><------><------><------># (change requires restart)
#unix_socket_group = "><------><------># (change requires restart)
#unix_socket_permissions = 0777><------># begin with 0 to use octal notation
<------><------><------><------><------># (change requires restart)
#bonjour = off<><------><------><------># advertise server via Bonjour
<------><------><------><------><------># (change requires restart)
#bonjour_name = "<----><------><------># defaults to the computer name
<------><------><------><------><------># (change requires restart)

# - Security and Authentication -

#authentication_timeout = 1min<><------># 1s-600s
ssl = true<----><------><------><------># (change requires restart)
#ssl_ciphers = 'HIGH:MEDIUM:+3DES:!aNULL' # allowed SSL ciphers
<------><------><------><------><------># (change requires restart)
#ssl_prefer_server_ciphers = on><------># (change requires restart)
#ssl_ecdh_curve = 'prime256v1'<><------># (change requires restart)
ssl_cert_file = '/etc/ssl/certs/ssl-cert-snakeoil.pem'<><------># (change requires restart)
ssl_key_file = '/etc/ssl/private/ssl-cert-snakeoil.key'><------># (change requires restart)
```

```
#ssl_ca_file = "<-----><------><------># (change requires restart)
#ssl_crl_file = "<----><------><------># (change requires restart)
#password_encryption = on
#db_user_namespace = off
#row_security = on

# GSSAPI using Kerberos
#krb_server_keyfile = ''
#krb_caseins_users = off

# - TCP Keepalives -
# see "man 7 tcp" for details

#tcp_keepalives_idle = 0<------><------># TCP_KEEPIDLE, in seconds;
<------><------><------><------><------># 0 selects the system default
#tcp_keepalives_interval = 0<--><------># TCP_KEEPINTVL, in seconds;
<------><------><------><------><------># 0 selects the system default
#tcp_keepalives_count = 0<-----><------># TCP_KEEPCNT;
<------><------><------><------><------># 0 selects the system default


#------------------------------------------------------------------------------
# RESOURCE USAGE (except WAL)
#------------------------------------------------------------------------------

# - Memory -

shared_buffers = 128MB<><------><------># min 128kB
<------><------><------><------><------># (change requires restart)
#huge_pages = try<-----><------><------># on, off, or try
<------><------><------><------><------># (change requires restart)
#temp_buffers = 8MB<---><------><------># min 800kB
#max_prepared_transactions = 0<><------># zero disables the feature
<------><------><------><------><------># (change requires restart)
# Caution: it is not advisable to set max_prepared_transactions nonzero unless
# you actively intend to use prepared transactions.
#work_mem = 4MB><------><------><------># min 64kB
#maintenance_work_mem = 64MB<--><------># min 1MB
#autovacuum_work_mem = -1<-----><------># min 1MB, or -1 to use maintenance_work_mem
#max_stack_depth = 2MB<><------><------># min 100kB
dynamic_shared_memory_type = posix<---->#  the default is the first option
<------><------><------><------><------># supported by the operating system:
<------><------><------><------><------>#   posix
<------><------><------><------><------>#   sysv
<------><------><------><------><------>#   windows
<------><------><------><------><------>#   mmap
<------><------><------><------><------># use none to disable dynamic shared memory

# - Disk -

#temp_file_limit = -1<-><------><------># limits per-session temp file space
<------><------><------><------><------># in kB, or -1 for no limit

# - Kernel Resource Usage -

#max_files_per_process = 1000<-><------># min 25
<------><------><------><------><------># (change requires restart)
#shared_preload_libraries = ''<><------># (change requires restart)

# - Cost-Based Vacuum Delay -

#vacuum_cost_delay = 0<><------><------># 0-100 milliseconds
#vacuum_cost_page_hit = 1<-----><------># 0-10000 credits
#vacuum_cost_page_miss = 10<---><------># 0-10000 credits
#vacuum_cost_page_dirty = 20<--><------># 0-10000 credits
#vacuum_cost_limit = 200<------><------># 1-10000 credits

# - Background Writer -

#bgwriter_delay = 200ms><------><------># 10-10000ms between rounds
#bgwriter_lru_maxpages = 100<--><------># 0-1000 max buffers written/round
#bgwriter_lru_multiplier = 2.0<><------># 0-10.0 multipler on buffers scanned/round
```

```
# - Asynchronous Behavior -

#effective_io_concurrency = 1<-><------># 1-1000; 0 disables prefetching
#max_worker_processes = 8


#------------------------------------------------------------------------------
# WRITE AHEAD LOG
#------------------------------------------------------------------------------

# - Settings -

#wal_level = minimal<--><------><------># minimal, archive, hot_standby, or logical
<------><------><------><------><------># (change requires restart)
#fsync = on<---><------><------><------># turns forced synchronization on or off
#synchronous_commit = on<------><------># synchronization level;
<------><------><------><------><------># off, local, remote_write, or on
#wal_sync_method = fsync<------><------># the default is the first option
<------><------><------><------><------># supported by the operating system:
<------><------><------><------><------>#   open_datasync
<------><------><------><------><------>#   fdatasync (default on Linux)
<------><------><------><------><------>#   fsync
<------><------><------><------><------>#   fsync_writethrough
<------><------><------><------><------>#   open_sync
#full_page_writes = on<><------><------># recover from partial page writes
#wal_compression = off<><------><------># enable compression of full-page writes
#wal_log_hints = off<--><------><------># also do full page writes of non-critical updates
<------><------><------><------><------># (change requires restart)
#wal_buffers = -1<-----><------><------># min 32kB, -1 sets based on shared_buffers
<------><------><------><------><------># (change requires restart)
#wal_writer_delay = 200ms<-----><------># 1-10000 milliseconds

#commit_delay = 0<-----><------><------># range 0-100000, in microseconds
#commit_siblings = 5<--><------><------># range 1-1000

# - Checkpoints -

#checkpoint_timeout = 5min<----><------># range 30s-1h
#max_wal_size = 1GB
#min_wal_size = 80MB
#checkpoint_completion_target = 0.5<---># checkpoint target duration, 0.0 - 1.0
#checkpoint_warning = 30s<-----><------># 0 disables

# - Archiving -

#archive_mode = off<---><------># enables archiving; off, on, or always
<------><------><------><------># (change requires restart)
#archive_command = "<-><------># command to use to archive a logfile segment
<------><------><------><------># placeholders: %p = path of file to archive
<------><------><------><------>#               %f = file name only
<------><------><------><------># e.g. 'test ! -f /mnt/server/archivedir/%f && cp %p
/mnt/server/archivedir/%f'
#archive_timeout = 0<--><------># force a logfile segment switch after this
<------><------><------><------># number of seconds; 0 disables


#------------------------------------------------------------------------------
# REPLICATION
#------------------------------------------------------------------------------

# - Sending Server(s) -

# Set these on the master and on any standby that will send replication data.

#max_wal_senders = 0<--><------># max number of walsender processes
<------><------><------><------># (change requires restart)
#wal_keep_segments = 0<><------># in logfile segments, 16MB each; 0 disables
#wal_sender_timeout = 60s<-----># in milliseconds; 0 disables

#max_replication_slots = 0<----># max number of replication slots
<------><------><------><------># (change requires restart)
#track_commit_timestamp = off<-># collect timestamp of transaction commit
<------><------><------><------># (change requires restart)
```

```
# - Master Server -

# These settings are ignored on a standby server.

#synchronous_standby_names = ''># standby servers that provide sync rep
<------><------><------><------># comma-separated list of application_name
<------><------><------><------># from standby(s); '*' = all
#vacuum_defer_cleanup_age = 0<->
# number of xacts by which cleanup is delayed

# - Standby Servers -

# These settings are ignored on a master server.

#hot_standby = off<----><------><------># "on" allows queries during recovery
<------><------><------><------><------># (change requires restart)
#max_standby_archive_delay = 30s<------># max delay before canceling queries
<------><------><------><------><------># when reading WAL from archive;
<------><------><------><------><------># -1 allows indefinite delay
#max_standby_streaming_delay = 30s<----># max delay before canceling queries
<------><------><------><------><------># when reading streaming WAL;
<------><------><------><------><------># -1 allows indefinite delay
#wal_receiver_status_interval = 10s<---># send replies at least this often
<------><------><------><------><------># 0 disables
#hot_standby_feedback = off<---><------># send info from standby to prevent
<------><------><------><------><------># query conflicts
#wal_receiver_timeout = 60s<---><------># time that receiver waits for
<------><------><------><------><------># communication from master
<------><------><------><------><------># in milliseconds; 0 disables
#wal_retrieve_retry_interval = 5s<-----># time to wait before retrying to
<------><------><------><------><------># retrieve WAL after a failed attempt


#------------------------------------------------------------------------------
# QUERY TUNING
#------------------------------------------------------------------------------

# - Planner Method Configuration -

#enable_bitmapscan = on
#enable_hashagg = on
#enable_hashjoin = on
#enable_indexscan = on
#enable_indexonlyscan = on
#enable_material = on
#enable_mergejoin = on
#enable_nestloop = on
#enable_seqscan = on
#enable_sort = on
#enable_tidscan = on

# - Planner Cost Constants -

#seq_page_cost = 1.0<--><------><------># measured on an arbitrary scale
#random_page_cost = 4.0><------><------># same scale as above
#cpu_tuple_cost = 0.01<><------><------># same scale as above
#cpu_index_tuple_cost = 0.005<-><------># same scale as above
#cpu_operator_cost = 0.0025<---><------># same scale as above
#effective_cache_size = 4GB

# - Genetic Query Optimizer -

#geqo = on
#geqo_threshold = 12
#geqo_effort = 5<------><------><------># range 1-10
#geqo_pool_size = 0<---><------><------># selects default based on effort
#geqo_generations = 0<-><------><------># selects default based on effort
#geqo_selection_bias = 2.0<----><------># range 1.5-2.0
#geqo_seed = 0.0<------><------><------># range 0.0-1.0

# - Other Planner Options -

#default_statistics_target = 100<------># range 1-10000
```

```
#constraint_exclusion = partition<----->
#cursor_tuple_fraction = 0.1<--><------>
#from_collapse_limit = 8
#join_collapse_limit = 8<------><------>
<------><------><------><------><------>

#------------------------------------------------------------------------------
# ERROR REPORTING AND LOGGING
#------------------------------------------------------------------------------

# - Where to Log -

#log_destination = 'stderr'<---><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>

# This is used when logging to stderr:
#logging_collector = off<------><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>

# These are only used if logging_collector is on:
#log_directory = 'pg_log'<-----><------>
<------><------><------><------><------>
#log_filename = 'postgresql-%Y-%m-%d_%H%M%S.log'<------>
<------><------><------><------><------>
#log_file_mode = 0600<-><------><------>
<------><------><------><------><------>
#log_truncate_on_rotation = off><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
#log_rotation_age = 1d<><------><------>
<------><------><------><------><------>
#log_rotation_size = 10MB<-----><------>
<------><------><------><------><------>
<------><------><------><------><------>

# These are relevant when logging to syslog:
#syslog_facility = 'LOCAL0'
#syslog_ident = 'postgres'

# This is only relevant when logging to eventlog (win32):
#event_source = 'PostgreSQL'

# - When to Log -

#client_min_messages = notice<-><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>

#log_min_messages = warning<---><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
<------><------><------><------><------>
```

```
<------><------><------><------><------>#   warning
<------><------><------><------><------>#   error
<------><------><------><------><------>#   log
<------><------><------><------><------>#   fatal
<------><------><------><------><------>#   panic

#log_min_error_statement = error<------># values in order of decreasing detail:
<------><------><------><------><------>#   debug5
<------><------><------><------><------>#   debug4
<------><------><------><------><------>#   debug3
<------><------><------><------><------>#   debug2
<------><------><------><------><------>#   debug1
<------><------><------><------><------>#   info
<------><------><------><------><------>#   notice
<------><------><------><------><------>#   warning
<------><------><------><------><------>#   error
<------><------><------><------><------>#   log
<------><------><------><------><------>#   fatal
<------><------><------><------><------>#   panic (effectively off)

#log_min_duration_statement = -1<------># -1 is disabled, 0 logs all statements
<------><------><------><------><------># and their durations, > 0 logs only
<------><------><------><------><------># statements running at least this number
<------><------><------><------><------># of milliseconds


# - What to Log -

#debug_print_parse = off
#debug_print_rewritten = off
#debug_print_plan = off
#debug_pretty_print = on
#log_checkpoints = off
#log_connections = off
#log_disconnections = off
#log_duration = off
#log_error_verbosity = default<><------># terse, default, or verbose messages
#log_hostname = off
log_line_prefix = '%t [%p-%l] %q%u@%d '><------><------># special values:
<------><------><------><------><------>#   %a = application name
<------><------><------><------><------>#   %u = user name
<------><------><------><------><------>#   %d = database name
<------><------><------><------><------>#   %r = remote host and port
<------><------><------><------><------>#   %h = remote host
<------><------><------><------><------>#   %p = process ID<------><------><------><------><------>#
%t = timestamp without milliseconds
<------><------><------><------><------>#   %m = timestamp with milliseconds
<------><------><------><------><------>#   %i = command tag
<------><------><------><------><------>#   %e = SQL state
<------><------><------><------><------>#   %c = session ID
<------><------><------><------><------>#   %l = session line number
<------><------><------><------><------>#   %s = session start timestamp
<------><------><------><------><------>#   %v = virtual transaction ID
<------><------><------><------><------>#   %x = transaction ID (0 if none)
<------><------><------><------><------>#   %q = stop here in non-session
<------><------><------><------><------>#        processes
<------><------><------><------><------>#   %% = '%'
<------><------><------><------><------># e.g. '<%u%%%d> '
#log_lock_waits = off<-><------><------># log lock waits >= deadlock_timeout
#log_statement = 'none'><------><------># none, ddl, mod, all
#log_replication_commands = off
#log_temp_files = -1<--><------><------># log temporary files equal or larger
<------><------><------><------><------># than the specified size in kilobytes;
<------><------><------><------><------># -1 disables, 0 logs all temp files
log_timezone = 'UTC'


# - Process Title -

#cluster_name = "<----><------><------># added to process titles if nonempty
<------><------><------><------><------># (change requires restart)
#update_process_title = on
```

```
#------------------------------------------------------------------------------
# RUNTIME STATISTICS
#------------------------------------------------------------------------------

# - Query/Index Statistics Collector -

#track_activities = on
#track_counts = on
#track_io_timing = off
#track_functions = none><------><------># none, pl, all
#track_activity_query_size = 1024<----->#(change requires restart)
stats_temp_directory = '/var/run/postgresql/9.5-main.pg_stat_tmp'


# - Statistics Monitoring -

#log_parser_stats = off
#log_planner_stats = off
#log_executor_stats = off
#log_statement_stats = off


#------------------------------------------------------------------------------
# AUTOVACUUM PARAMETERS
#------------------------------------------------------------------------------

#autovacuum = on<------><------><------># Enable autovacuum subprocess?  'on'
<------><------><------><------><------># requires track_counts to also be on.
#log_autovacuum_min_duration = -1<----->#-1 disables, 0 logs all actions and
<------><------><------><------><------># their durations, > 0 logs only
<------><------><------><------><------># actions running at least this number
<------><------><------><------><------># of milliseconds.
#autovacuum_max_workers = 3<--->#<------># max number of autovacuum subprocesses
<------><------><------><------><------># (change requires restart)
#autovacuum_naptime = 1min<---><------># time between autovacuum runs
#autovacuum_vacuum_threshold = 50<----->#min number of row updates before
<------><------><------><------><------># vacuum
#autovacuum_analyze_threshold = 50<---->#min number of row updates before
<------><------><------><------><------># analyze
#autovacuum_vacuum_scale_factor = 0.2<->#fraction of table size before vacuum
#autovacuum_analyze_scale_factor = 0.1<>#fraction of table size before analyze
#autovacuum_freeze_max_age = 200000000<>#maximum XID age before forced vacuum
<------><------><------><------><------># (change requires restart)
#autovacuum_multixact_freeze_max_age = 400000000<------># maximum multixact age
<------><------><------><------><------># before forced vacuum
<------><------><------><------><------># (change requires restart)
#autovacuum_vacuum_cost_delay = 20ms<-->#default vacuum cost delay for
<------><------><------><------><------># autovacuum, in milliseconds;
<------><------><------><------><------># -1 means use vacuum_cost_delay
#autovacuum_vacuum_cost_limit = -1<---->#default vacuum cost limit for
<------><------><------><------><------># autovacuum, -1 means use
<------><------><------><------><------># vacuum_cost_limit


#------------------------------------------------------------------------------
# CLIENT CONNECTION DEFAULTS
#------------------------------------------------------------------------------

# - Statement Behavior -

#search_path = '"$user", public'<------># schema names
#default_tablespace = ''<------><------># a tablespace name, '' uses the default
#temp_tablespaces = ''<><------><------># a list of tablespace names, '' uses
<------><------><------><------><------># only default tablespace
#check_function_bodies = on
#default_transaction_isolation = 'read committed'
#default_transaction_read_only = off
#default_transaction_deferrable = off
#session_replication_role = 'origin'
#statement_timeout = 0<><------><------># in milliseconds, 0 is disabled
#lock_timeout = 0<-----><------><------># in milliseconds, 0 is disabled
#vacuum_freeze_min_age = 50000000
```

```
#vacuum_freeze_table_age = 150000000
#vacuum_multixact_freeze_min_age = 5000000
#vacuum_multixact_freeze_table_age = 150000000
#bytea_output = 'hex'<->><------><------># hex, escape
#xmlbinary = 'base64'
#xmloption = 'content'
#gin_fuzzy_search_limit = 0
#gin_pending_list_limit = 4MB

# - Locale and Formatting -

datestyle = 'iso, mdy'
#intervalstyle = 'postgres'
timezone = 'UTC'
#timezone_abbreviations = 'Default'     # Select the set of available time zone
<------><------><------><------><------># abbreviations.  Currently, there are
<------><------><------><------><------>#   Default
<------><------><------><------><------>#   Australia (historical usage)
<------><------><------><------><------>#   India
<------><------><------><------><------># You can create your own file in
<------><------><------><------><------># share/timezonesets/.
#extra_float_digits = 0><------><------># min -15, max 3
#client_encoding = sql_ascii<--><------># actually, defaults to database
<------><------><------><------><------># encoding

# These settings are initialized by initdb, but they can be changed.
lc_messages = 'en_US.UTF-8'<---><------><------># locale for system error message
<------><------><------><------><------># strings
lc_monetary = 'en_US.UTF-8'<---><------><------># locale for monetary formatting
lc_numeric = 'en_US.UTF-8'<----><------><------># locale for number formatting
lc_time = 'en_US.UTF-8'><------><------><------># locale for time formatting

# default configuration for text search
default_text_search_config = 'pg_catalog.english'

# - Other Defaults -

#dynamic_library_path = '$libdir'
#local_preload_libraries = ''
#session_preload_libraries = ''


#------------------------------------------------------------------------------
# LOCK MANAGEMENT
#------------------------------------------------------------------------------

#deadlock_timeout = 1s
#max_locks_per_transaction = 64><------># min 10
<------><------><------><------><------># (change requires restart)
#max_pred_locks_per_transaction = 64<--># min 10
<------><------><------><------><------># (change requires restart)


#------------------------------------------------------------------------------
# VERSION/PLATFORM COMPATIBILITY
#------------------------------------------------------------------------------

# - Previous PostgreSQL Versions -

#array_nulls = on
#backslash_quote = safe_encoding<------># on, off, or safe_encoding
#default_with_oids = off
#escape_string_warning = on
#lo_compat_privileges = off
#operator_precedence_warning = off
#quote_all_identifiers = off
#sql_inheritance = on
#standard_conforming_strings = on
#synchronize_seqscans = on

# - Other Platforms and Clients -

#transform_null_equals = off
```

```
#------------------------------------------------------------------------------
# ERROR HANDLING
#------------------------------------------------------------------------------

#exit_on_error = off<--><------><------># terminate session on any error?
#restart_after_crash = on<-----><------># reinitialize after backend crash?


#------------------------------------------------------------------------------
# CONFIG FILE INCLUDES
#------------------------------------------------------------------------------

# These options allow settings to be loaded from files other than the
# default postgresql.conf.

#include_dir = 'conf.d'><------><------># include files ending in '.conf' from
<------><------><------><------><------># directory 'conf.d'
#include_if_exists = 'exists.conf'<---->#include file only if it exists
#include = 'special.conf'<-----><------># include file


#------------------------------------------------------------------------------
# CUSTOMIZED OPTIONS
#------------------------------------------------------------------------------

# Add settings for extensions here
```

I. Nazwa serwera: orav1

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	orav1
IP / Domena	10.90.54.11
OS	Windows Server 2012 R2 Standard
RAM:	32 GB
vCPU	10
Przestrzeń	1900 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Otwarte Repozytorium A/V

2. Konfiguracja .hosts

10.90.54.31	rt.epo.pl
10.90.54.250	av.epo.pl
10.90.54.250	av.epodreczniki.pl

3. Interfejsy sieciowe

LAN 10.90.54.11

4. Zainstalowane usługi

a. Bazy danych

Parametr	Wartość
Lokalizacja instalacji	MSSQL Server 2014 Express MaterialFormatDB C:\Program Files\Microsoft SQL Server\MSSQL12.MSSQLSERVER\MSSQL\DATA
Lokalizacja instalacji	MSSQL Server 2014 Express RepDB C:\Program Files\Microsoft SQL Server\MSSQL12.MSSQLSERVER\MSSQL\DATA

b. RepoAV

Parametr	Wartość
Lokalizacja instalacji	C:\RepoAV E:\Storage
Zawartość pliku konfiguracyjnego #1	IIS - C:\RepoAV\Services\RepAPI\Web.config <?xml version="1.0" encoding="utf-8"?> <!-- For more information on how to configure your ASP.NET application, please visit http://go.microsoft.com/fwlink/?LinkId=169433 --> <configuration> <configSections> <section name="loggingConfiguration" type="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.LoggingSettings, Microsoft.Practices.EnterpriseLibrary.Logging" /> </configSections> <connectionStrings> <add name="DefaultConnection" providerName="System.Data.SqlClient" connectionString="Data Source=(local);Initial Catalog=RepDB;User ID=sa;Password=ore!infomex!2017" /> </connectionStrings> <appSettings> <add key="webpages:Version" value="2.0.0.0" /> <add key="webpages:Enabled" value="false" /> <add key="PreserveLoginUrl" value="true" /> <add key="ClientValidationEnabled" value="true" /> <add key="UnobtrusiveJavaScriptEnabled" value="true" /> </appSettings> <loggingConfiguration name="Logging Application Block" tracingEnabled="false" defaultCategory="" logWarningsWhenNoCategoriesMatch="false"> <listeners> <add fileName="C:\RepoAV\Logs\RepAPI.log" header="" footer="" formatter="Text Formatter" listenerDataType="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.FlatFileTraceListenerData, Microsoft.Practices.EnterpriseLibrary.Logging" traceOutputOptions="None" type="Microsoft.Practices.EnterpriseLibrary.Logging.TraceListeners.FlatFileTraceListener, Microsoft.Practices.EnterpriseLibrary.Logging" name="FlatFile TraceListener" /> </listeners> <formatters> <add template="[{timestamp(local:yyyy-MM-dd HH:mm:ss.ffff)}]
[{severity}][{category}] {message}
" type="Microsoft.Practices.EnterpriseLibrary.Logging.Formatters.TextFor

```
matter, Microsoft.Practices.EnterpriseLibrary.Logging" name="Text
Formatter" />
    </formatters>
    <categorySources>
      <add switchValue="Information" name="General">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </add>
      <add switchValue="Information" name="RepositoryAccess">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </add>
    </categorySources>
    <specialSources>
      <allEvents switchValue="All" name="All Events" />
      <notProcessed switchValue="All" name="Unprocessed Category" />
      <errors switchValue="All" name="Logging Errors &amp;
Warnings">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </errors>
    </specialSources>
  </loggingConfiguration>
  <system.web>
    <compilation debug="true" targetFramework="4.5" />
    <httpRuntime targetFramework="4.5" />
    <authentication mode="None" />
  </system.web>
  <system.webServer>
    <validation validateIntegratedModeConfiguration="false" />
    <handlers>
      <remove name="ExtensionlessUrlHandler-Integrated-4.0" />
      <remove name="OPTIONSVerbHandler" />
      <remove name="TRACEVerbHandler" />
      <add name="ExtensionlessUrlHandler-Integrated-4.0" path="*."
verb="*" type="System.Web.Handlers.TransferRequestHandler"
preCondition="integratedMode,runtimeVersionv4.0" />
    </handlers>
        <httpProtocol>
          <customHeaders>
            <remove name="Access-Control-Allow-Origin" />
            <add name="Access-Control-Allow-Origin" value="*" />
          </customHeaders>
        </httpProtocol>
  </system.webServer>
```

	<runtime> <assemblyBinding xmlns="urn:schemas-microsoft-com:asm.v1"> <dependentAssembly> <assemblyIdentity name="System.Web.Helpers" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-2.0.0.0" newVersion="2.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="System.Web.Mvc" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-4.0.0.0" newVersion="4.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="System.Web.WebPages" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-2.0.0.0" newVersion="2.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="WebGrease" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="0.0.0.0-1.3.0.0" newVersion="1.3.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="Newtonsoft.Json" publicKeyToken="30ad4fe6b2a6aeed" culture="neutral" /> <bindingRedirect oldVersion="0.0.0.0-6.0.0.0" newVersion="6.0.0.0" /> </dependentAssembly> </assemblyBinding> </runtime> </configuration>
Zawartość pliku konfiguracyjnego #2	IIS - C:\RepoAV\Services\RepositoryAccess\Web.config <?xml version="1.0" encoding="UTF-8"?> <!-- For more information on how to configure your ASP.NET application, please visit http://go.microsoft.com/fwlink/?LinkId=169433 --> <configuration> <configSections> <section name="loggingConfiguration" type="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.LoggingSettings, Microsoft.Practices.EnterpriseLibrary.Logging" />

```
    <section name="repositoryNodesConfiguration"
type="PSNC.RepoAV.Services.RepositoryAccess.RepositoryNodesConfi
gurationSection, RepositoryAccess" />
  </configSections>
  <repositoryNodesConfiguration>
    <redirection enabled="true" smartRedirectForRequestHeader="EP-
CDN-REQ" />
    <repositoryNodes thisNodeId="1">
      <add id="1" enabled="true" address="10.90.54.11" />
      <add id="2" enabled="true" address="10.90.54.12" />
      <add id="3" enabled="true" address="10.90.54.13" />
            <add id="4" enabled="true" address="10.90.54.14" />
            <add id="5" enabled="true" address="10.90.54.15" />
            <add id="6" enabled="true" address="10.90.54.16" />
    </repositoryNodes>
  </repositoryNodesConfiguration>
  <appSettings>
    <add key="NetworkAddressAndMaskWithoutAuthorization"
value="10.90.54.0/24" />
    <add key="ResponseCachingTime" value="5" />
  </appSettings>
  <connectionStrings>
    <add name="ContentDB" connectionString="Data
Source=(local);Initial Catalog=MaterialFormatDB;User
ID=sa;Password=ore!infomex!2017"
providerName="System.Data.SqlClient" />
    <add name="RepoDB" connectionString="Data Source=(local);Initial
Catalog=RepDB;User ID=sa;Password=ore!infomex!2017"
providerName="System.Data.SqlClient" />
  </connectionStrings>
  <loggingConfiguration name="Logging Application Block"
tracingEnabled="false" defaultCategory=""
logWarningsWhenNoCategoriesMatch="false">
    <listeners>
      <add fileName="C:\RepoAV\Logs\RepositoryAccess2.log" header=""
footer="" formatter="Text Formatter"
listenerDataType="Microsoft.Practices.EnterpriseLibrary.Logging.Config
uration.FlatFileTraceListenerData,
Microsoft.Practices.EnterpriseLibrary.Logging"
traceOutputOptions="None"
type="Microsoft.Practices.EnterpriseLibrary.Logging.TraceListeners.Flat
FileTraceListener, Microsoft.Practices.EnterpriseLibrary.Logging"
name="FlatFile TraceListener" />
    </listeners>
    <formatters>
      <add template="[{timestamp(local:yyyy-MM-dd
HH:mm:ss.ffff)}]&#xA;[{severity}][{category}] {message}&#xA;"
type="Microsoft.Practices.EnterpriseLibrary.Logging.Formatters.TextFor
```

```
matter, Microsoft.Practices.EnterpriseLibrary.Logging" name="Text
Formatter" />
    </formatters>
    <categorySources>
     <add switchValue="Information" name="General">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </add>
     <add switchValue="Information" name="RepositoryAccess">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </add>
    </categorySources>
    <specialSources>
     <allEvents switchValue="All" name="All Events" />
     <notProcessed switchValue="All" name="Unprocessed Category" />
     <errors switchValue="All" name="Logging Errors &amp;
Warnings">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </errors>
    </specialSources>
  </loggingConfiguration>
  <system.web>
    <compilation debug="true" targetFramework="4.5" />
    <httpRuntime targetFramework="4.5" />
  </system.web>
  <system.webServer>
    <modules>
     <remove name="DownloadModule" />
     <add name="DownloadModule"
type="PSNC.RepoAV.Services.RepositoryAccess.DownloadModule" />
    </modules>
      <httpProtocol>
        <customHeaders>
       <add name="Access-Control-Allow-Headers" value="X-
Requested-With" />
        </customHeaders>
      </httpProtocol>
  </system.webServer>
</configuration>
```

I. Nazwa serwera: orav2

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	orav2
IP / Domena	10.90.54.12
OS	Windows Server 2012 R2 Standard
RAM:	32 GB
vCPU	10
Przestrzeń	1850 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Otwarte Repozytorium A/V

2. Konfiguracja .hosts

10.90.54.31	rt.epo.pl
10.90.54.250	av.epo.pl
10.90.54.250	av.epodreczniki.pl

3. Interfejsy sieciowe

LAN 10.90.54.12

4. Zainstalowane usługi

a. Bazy danych

Parametr	Wartość
Lokalizacja instalacji	MSSQL Server 2014 Express MaterialFormatDB C:\Program Files\Microsoft SQL Server\MSSQL12.MSSQLSERVER\MSSQL\DATA

b. RepoAV

Parametr	Wartość
Lokalizacja instalacji	C:\RepoAV E:\Storage
Zawartość pliku konfiguracyjnego #1	IIS - C:\RepoAV\Services\RepAPI\Web.config <?xml version="1.0" encoding="utf-8"?> <!-- For more information on how to configure your ASP.NET application, please visit http://go.microsoft.com/fwlink/?LinkId=169433 --> <configuration> <configSections> <section name="loggingConfiguration" type="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.LoggingSettings, Microsoft.Practices.EnterpriseLibrary.Logging" /> </configSections> <connectionStrings> <add name="DefaultConnection" providerName="System.Data.SqlClient" connectionString="Data Source=ORAV1;Initial Catalog=RepDB;User ID=sa;Password=ore!infomex!2017" /> </connectionStrings> <appSettings> <add key="webpages:Version" value="2.0.0.0" /> <add key="webpages:Enabled" value="false" /> <add key="PreserveLoginUrl" value="true" /> <add key="ClientValidationEnabled" value="true" /> <add key="UnobtrusiveJavaScriptEnabled" value="true" /> </appSettings> <loggingConfiguration name="Logging Application Block" tracingEnabled="false" defaultCategory="" logWarningsWhenNoCategoriesMatch="false"> <listeners> <add fileName="C:\RepoAV\Logs\RepAPI.log" header="" footer="" formatter="Text Formatter" listenerDataType="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.FlatFileTraceListenerData, Microsoft.Practices.EnterpriseLibrary.Logging" traceOutputOptions="None" type="Microsoft.Practices.EnterpriseLibrary.Logging.TraceListeners.FlatFileTraceListener, Microsoft.Practices.EnterpriseLibrary.Logging" name="FlatFile TraceListener" /> </listeners> <formatters> <add template="[{timestamp(local:yyyy-MM-dd HH:mm:ss.ffff)}]
[{severity}][{category}] {message}
" type="Microsoft.Practices.EnterpriseLibrary.Logging.Formatters.TextFor

```
matter, Microsoft.Practices.EnterpriseLibrary.Logging" name="Text
Formatter" />
    </formatters>
    <categorySources>
      <add switchValue="Information" name="General">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </add>
      <add switchValue="Information" name="RepositoryAccess">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </add>
    </categorySources>
    <specialSources>
      <allEvents switchValue="All" name="All Events" />
      <notProcessed switchValue="All" name="Unprocessed Category" />
      <errors switchValue="All" name="Logging Errors &amp;
Warnings">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </errors>
    </specialSources>
  </loggingConfiguration>
  <system.web>
    <compilation debug="true" targetFramework="4.5" />
    <httpRuntime targetFramework="4.5" />
    <authentication mode="None" />
  </system.web>
  <system.webServer>
    <validation validateIntegratedModeConfiguration="false" />
    <handlers>
      <remove name="ExtensionlessUrlHandler-Integrated-4.0" />
      <remove name="OPTIONSVerbHandler" />
      <remove name="TRACEVerbHandler" />
      <add name="ExtensionlessUrlHandler-Integrated-4.0" path="*."
verb="*" type="System.Web.Handlers.TransferRequestHandler"
preCondition="integratedMode,runtimeVersionv4.0" />
    </handlers>
        <httpProtocol>
          <customHeaders>
            <remove name="Access-Control-Allow-Headers" />
            <add name="Access-Control-Allow-Origin" value="*" />
          </customHeaders>
        </httpProtocol>
  </system.webServer>
```

	<runtime> <assemblyBinding xmlns="urn:schemas-microsoft-com:asm.v1"> <dependentAssembly> <assemblyIdentity name="System.Web.Helpers" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-2.0.0.0" newVersion="2.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="System.Web.Mvc" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-4.0.0.0" newVersion="4.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="System.Web.WebPages" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-2.0.0.0" newVersion="2.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="WebGrease" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="0.0.0.0-1.3.0.0" newVersion="1.3.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="Newtonsoft.Json" publicKeyToken="30ad4fe6b2a6aeed" culture="neutral" /> <bindingRedirect oldVersion="0.0.0.0-6.0.0.0" newVersion="6.0.0.0" /> </dependentAssembly> </assemblyBinding> </runtime> </configuration>
Zawartość pliku konfiguracyjnego #2	IIS - C:\RepoAV\Services\RepositoryAccess\Web.config <?xml version="1.0" encoding="UTF-8"?> <!-- For more information on how to configure your ASP.NET application, please visit http://go.microsoft.com/fwlink/?LinkId=169433 --> <configuration> <configSections> <section name="loggingConfiguration" type="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.LoggingSettings, Microsoft.Practices.EnterpriseLibrary.Logging" />

```
    <section name="repositoryNodesConfiguration"
type="PSNC.RepoAV.Services.RepositoryAccess.RepositoryNodesConfi
gurationSection, RepositoryAccess" />
  </configSections>
  <repositoryNodesConfiguration>
    <redirection enabled="true" smartRedirectForRequestHeader="EP-
CDN-REQ" />
    <repositoryNodes thisNodeId="2">
      <add id="2" enabled="true" address="10.90.54.12" />
      <add id="1" enabled="true" address="10.90.54.11" />
      <add id="3" enabled="true" address="10.90.54.13" />
            <add id="4" enabled="true" address="10.90.54.14" />
            <add id="5" enabled="true" address="10.90.54.15" />
            <add id="6" enabled="true" address="10.90.54.16" />
    </repositoryNodes>
  </repositoryNodesConfiguration>
  <appSettings>
    <add key="NetworkAddressAndMaskWithoutAuthorization"
value="10.90.54.0/24" />
    <add key="ResponseCachingTime" value="5" />
  </appSettings>
  <connectionStrings>
    <add name="ContentDB" connectionString="Data
Source=(local);Initial Catalog=MaterialFormatDB;Integrated
Security=SSPI" providerName="System.Data.SqlClient" />
    <add name="RepoDB" connectionString="Data Source=ORAV1;Initial
Catalog=RepDB;User ID=sa;Password=ore!infomex!2017"
providerName="System.Data.SqlClient" />
  </connectionStrings>
  <loggingConfiguration name="Logging Application Block"
tracingEnabled="false" defaultCategory=""
logWarningsWhenNoCategoriesMatch="false">
    <listeners>
      <add fileName="C:\RepoAV\Logs\RepositoryAccess2.log" header=""
footer="" formatter="Text Formatter"
listenerDataType="Microsoft.Practices.EnterpriseLibrary.Logging.Config
uration.FlatFileTraceListenerData,
Microsoft.Practices.EnterpriseLibrary.Logging"
traceOutputOptions="None"
type="Microsoft.Practices.EnterpriseLibrary.Logging.TraceListeners.Flat
FileTraceListener, Microsoft.Practices.EnterpriseLibrary.Logging"
name="FlatFile TraceListener" />
    </listeners>
    <formatters>
      <add template="[{timestamp(local:yyyy-MM-dd
HH:mm:ss.ffff)}]&#xA;[{severity}][{category}] {message}&#xA;"
type="Microsoft.Practices.EnterpriseLibrary.Logging.Formatters.TextFor
```

```
matter, Microsoft.Practices.EnterpriseLibrary.Logging" name="Text
Formatter" />
    </formatters>
    <categorySources>
     <add switchValue="Information" name="General">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </add>
     <add switchValue="Information" name="RepositoryAccess">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </add>
    </categorySources>
    <specialSources>
     <allEvents switchValue="All" name="All Events" />
     <notProcessed switchValue="All" name="Unprocessed Category" />
     <errors switchValue="All" name="Logging Errors &amp;
Warnings">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </errors>
    </specialSources>
  </loggingConfiguration>
  <system.web>
    <compilation debug="true" targetFramework="4.5" />
    <httpRuntime targetFramework="4.5" />
  </system.web>
  <system.webServer>
    <modules>
     <remove name="DownloadModule" />
     <add name="DownloadModule"
type="PSNC.RepoAV.Services.RepositoryAccess.DownloadModule" />
    </modules>
      <httpProtocol>
        <customHeaders>
    <add name="Access-Control-Allow-Headers" value="X-Requested-
With" />
        </customHeaders>
      </httpProtocol>
  </system.webServer>
</configuration>
```

I. Nazwa serwera: orav3

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	orav3
IP / Domena	10.90.54.13
OS	Windows Server 2012 R2 Standard
RAM:	32 GB
vCPU	10
Przestrzeń	800 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Otwarte Repozytorium A/V

2. Konfiguracja .hosts

10.90.54.31	rt.epo.pl
10.90.54.250	av.epo.pl
10.90.54.250	av.epodreczniki.pl

3. Interfejsy sieciowe

LAN 10.90.54.13

4. Zainstalowane usługi

a. Bazy danych

Parametr	Wartość
Lokalizacja instalacji	MSSQL Server 2014 Express MaterialFormatDB C:\Program Files\Microsoft SQL Server\MSSQL12.MSSQLSERVER\MSSQL\DATA

b. RepoAV

Parametr	Wartość
Lokalizacja instalacji	C:\RepoAV E:\Storage
Zawartość pliku konfiguracyjnego #1	IIS - C:\RepoAV\Services\RepAPI\Web.config <?xml version="1.0" encoding="utf-8"?> <!-- For more information on how to configure your ASP.NET application, please visit http://go.microsoft.com/fwlink/?LinkId=169433 --> <configuration> <configSections> <section name="loggingConfiguration" type="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.LoggingSettings, Microsoft.Practices.EnterpriseLibrary.Logging" /> </configSections> <connectionStrings> <add name="DefaultConnection" providerName="System.Data.SqlClient" connectionString="Data Source=ORAV1;Initial Catalog=RepDB;User ID=sa;Password=ore!infomex!2017" /> </connectionStrings> <appSettings> <add key="webpages:Version" value="2.0.0.0" /> <add key="webpages:Enabled" value="false" /> <add key="PreserveLoginUrl" value="true" /> <add key="ClientValidationEnabled" value="true" /> <add key="UnobtrusiveJavaScriptEnabled" value="true" /> </appSettings> <loggingConfiguration name="Logging Application Block" tracingEnabled="false" defaultCategory="" logWarningsWhenNoCategoriesMatch="false"> <listeners> <add fileName="C:\RepoAV\Logs\RepAPI.log" header="" footer="" formatter="Text Formatter" listenerDataType="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.FlatFileTraceListenerData, Microsoft.Practices.EnterpriseLibrary.Logging" traceOutputOptions="None" type="Microsoft.Practices.EnterpriseLibrary.Logging.TraceListeners.FlatFileTraceListener, Microsoft.Practices.EnterpriseLibrary.Logging" name="FlatFile TraceListener" /> </listeners> <formatters> <add template="[{timestamp(local:yyyy-MM-dd HH:mm:ss.ffff)}]
[{severity}][{category}] {message}
" type="Microsoft.Practices.EnterpriseLibrary.Logging.Formatters.TextFor

```
matter, Microsoft.Practices.EnterpriseLibrary.Logging" name="Text
Formatter" />
    </formatters>
    <categorySources>
      <add switchValue="Information" name="General">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </add>
      <add switchValue="Information" name="RepositoryAccess">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </add>
    </categorySources>
    <specialSources>
      <allEvents switchValue="All" name="All Events" />
      <notProcessed switchValue="All" name="Unprocessed Category" />
      <errors switchValue="All" name="Logging Errors &amp;
Warnings">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </errors>
    </specialSources>
  </loggingConfiguration>
  <system.web>
    <compilation debug="true" targetFramework="4.5" />
    <httpRuntime targetFramework="4.5" />
    <authentication mode="None" />
  </system.web>
  <system.webServer>
    <validation validateIntegratedModeConfiguration="false" />
    <handlers>
      <remove name="ExtensionlessUrlHandler-Integrated-4.0" />
      <remove name="OPTIONSVerbHandler" />
      <remove name="TRACEVerbHandler" />
      <add name="ExtensionlessUrlHandler-Integrated-4.0" path="*."
verb="*" type="System.Web.Handlers.TransferRequestHandler"
preCondition="integratedMode,runtimeVersionv4.0" />
    </handlers>
        <httpProtocol>
          <customHeaders>
            <add name="Access-Control-Allow-Origin" value="*" />
          </customHeaders>
        </httpProtocol>
  </system.webServer>
  <runtime>
```

	<assemblyBinding xmlns="urn:schemas-microsoft-com:asm.v1"> <dependentAssembly> <assemblyIdentity name="System.Web.Helpers" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-2.0.0.0" newVersion="2.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="System.Web.Mvc" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-4.0.0.0" newVersion="4.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="System.Web.WebPages" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-2.0.0.0" newVersion="2.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="WebGrease" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="0.0.0.0-1.3.0.0" newVersion="1.3.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="Newtonsoft.Json" publicKeyToken="30ad4fe6b2a6aeed" culture="neutral" /> <bindingRedirect oldVersion="0.0.0.0-6.0.0.0" newVersion="6.0.0.0" /> </dependentAssembly> </assemblyBinding> </runtime> </configuration>
Zawartość pliku konfiguracyjnego #2	IIS - C:\RepoAV\Services\RepositoryAccess\Web.config <?xml version="1.0" encoding="UTF-8"?> <!-- For more information on how to configure your ASP.NET application, please visit http://go.microsoft.com/fwlink/?LinkId=169433 --> <configuration> <configSections> <section name="loggingConfiguration" type="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.LoggingSettings, Microsoft.Practices.EnterpriseLibrary.Logging" />

```
    <section name="repositoryNodesConfiguration"
type="PSNC.RepoAV.Services.RepositoryAccess.RepositoryNodesConfi
gurationSection, RepositoryAccess" />
  </configSections>
  <repositoryNodesConfiguration>
    <redirection enabled="true" smartRedirectForRequestHeader="EP-
CDN-REQ" />
    <repositoryNodes thisNodeId="3">
      <add id="3" enabled="true" address="10.90.54.13" />
      <add id="1" enabled="true" address="10.90.54.11" />
      <add id="2" enabled="true" address="10.90.54.12" />
            <add id="4" enabled="true" address="10.90.54.14" />
            <add id="5" enabled="true" address="10.90.54.15" />
            <add id="6" enabled="true" address="10.90.54.16" />
    </repositoryNodes>
  </repositoryNodesConfiguration>
  <appSettings>
    <add key="NetworkAddressAndMaskWithoutAuthorization"
value="10.90.54.0/24" />
    <add key="ResponseCachingTime" value="5" />
  </appSettings>
  <connectionStrings>
    <add name="ContentDB" connectionString="Data
Source=(local);Initial Catalog=MaterialFormatDB;Integrated
Security=SSPI" providerName="System.Data.SqlClient" />
    <add name="RepoDB" connectionString="Data Source=ORAV1;Initial
Catalog=RepDB;User ID=sa;Password=ore!infomex!2017"
providerName="System.Data.SqlClient" />
  </connectionStrings>
  <loggingConfiguration name="Logging Application Block"
tracingEnabled="false" defaultCategory=""
logWarningsWhenNoCategoriesMatch="false">
    <listeners>
      <add fileName="C:\RepoAV\Logs\RepositoryAccess2.log" header=""
footer="" formatter="Text Formatter"
listenerDataType="Microsoft.Practices.EnterpriseLibrary.Logging.Config
uration.FlatFileTraceListenerData,
Microsoft.Practices.EnterpriseLibrary.Logging"
traceOutputOptions="None"
type="Microsoft.Practices.EnterpriseLibrary.Logging.TraceListeners.Flat
FileTraceListener, Microsoft.Practices.EnterpriseLibrary.Logging"
name="FlatFile TraceListener" />
    </listeners>
    <formatters>
      <add template="[{timestamp(local:yyyy-MM-dd
HH:mm:ss.ffff)}]&#xA;[{severity}][{category}] {message}&#xA;"
type="Microsoft.Practices.EnterpriseLibrary.Logging.Formatters.TextFor
```

```
matter, Microsoft.Practices.EnterpriseLibrary.Logging" name="Text
Formatter" />
    </formatters>
    <categorySources>
     <add switchValue="Information" name="General">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </add>
     <add switchValue="Information" name="RepositoryAccess">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </add>
    </categorySources>
    <specialSources>
     <allEvents switchValue="All" name="All Events" />
     <notProcessed switchValue="All" name="Unprocessed Category" />
     <errors switchValue="All" name="Logging Errors &amp;
Warnings">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </errors>
    </specialSources>
  </loggingConfiguration>
  <system.web>
    <compilation debug="true" targetFramework="4.5" />
    <httpRuntime targetFramework="4.5" />
  </system.web>
  <system.webServer>
    <modules>
     <remove name="DownloadModule" />
     <add name="DownloadModule"
type="PSNC.RepoAV.Services.RepositoryAccess.DownloadModule" />
    </modules>
      <httpProtocol>
        <customHeaders>
    <add name="Access-Control-Allow-Headers" value="X-Requested-
With" />
        </customHeaders>
      </httpProtocol>
  </system.webServer>
</configuration>
```

I. Nazwa serwera: orav4

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	orav4
IP / Domena	10.90.54.14
OS	Windows Server 2012 R2 Standard
RAM:	32 GB
vCPU	10
Przestrzeń	1000 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Otwarte Repozytorium A/V

2. Konfiguracja .hosts

10.90.54.31	rt.epo.pl
10.90.54.250	av.epo.pl
10.90.54.250	av.epodreczniki.pl

3. Interfejsy sieciowe

LAN 10.90.54.14

4. Zainstalowane usługi

a. Bazy danych

Parametr	Wartość
Lokalizacja instalacji	MSSQL Server 2014 Express MaterialFormatDB C:\Program Files\Microsoft SQL Server\MSSQL12.MSSQLSERVER\MSSQL\DATA

b. RepoAV

Parametr	Wartość
Lokalizacja instalacji	C:\RepoAV E:\Storage
Zawartość pliku konfiguracyjnego #1	IIS - C:\RepoAV\Services\RepAPI\Web.config <?xml version="1.0" encoding="utf-8"?> <!-- For more information on how to configure your ASP.NET application, please visit http://go.microsoft.com/fwlink/?LinkId=169433 --> <configuration> <configSections> <section name="loggingConfiguration" type="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.LoggingSettings, Microsoft.Practices.EnterpriseLibrary.Logging" /> </configSections> <connectionStrings> <add name="DefaultConnection" providerName="System.Data.SqlClient" connectionString="Data Source=ORAV1;Initial Catalog=RepDB;User ID=sa;Password=ore!infomex!2017" /> </connectionStrings> <appSettings> <add key="webpages:Version" value="2.0.0.0" /> <add key="webpages:Enabled" value="false" /> <add key="PreserveLoginUrl" value="true" /> <add key="ClientValidationEnabled" value="true" /> <add key="UnobtrusiveJavaScriptEnabled" value="true" /> </appSettings> <loggingConfiguration name="Logging Application Block" tracingEnabled="false" defaultCategory="" logWarningsWhenNoCategoriesMatch="false"> <listeners> <add fileName="C:\RepoAV\Logs\RepAPI.log" header="" footer="" formatter="Text Formatter" listenerDataType="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.FlatFileTraceListenerData, Microsoft.Practices.EnterpriseLibrary.Logging" traceOutputOptions="None" type="Microsoft.Practices.EnterpriseLibrary.Logging.TraceListeners.FlatFileTraceListener, Microsoft.Practices.EnterpriseLibrary.Logging" name="FlatFile TraceListener" /> </listeners> <formatters> <add template="[{timestamp(local:yyyy-MM-dd HH:mm:ss.ffff)}]
[{severity}][{category}] {message}
"

```
type="Microsoft.Practices.EnterpriseLibrary.Logging.Formatters.TextFormatter, Microsoft.Practices.EnterpriseLibrary.Logging" name="Text Formatter" />
    </formatters>
    <categorySources>
      <add switchValue="Information" name="General">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </add>
      <add switchValue="Information" name="RepositoryAccess">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </add>
    </categorySources>
    <specialSources>
      <allEvents switchValue="All" name="All Events" />
      <notProcessed switchValue="All" name="Unprocessed Category" />
      <errors switchValue="All" name="Logging Errors &amp;amp; Warnings">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </errors>
    </specialSources>
  </loggingConfiguration>
  <system.web>
    <compilation debug="true" targetFramework="4.5" />
    <httpRuntime targetFramework="4.5" />
    <authentication mode="None" />
  </system.web>
  <system.webServer>
    <validation validateIntegratedModeConfiguration="false" />
    <handlers>
      <remove name="ExtensionlessUrlHandler-Integrated-4.0" />
      <remove name="OPTIONSVerbHandler" />
      <remove name="TRACEVerbHandler" />
      <add name="ExtensionlessUrlHandler-Integrated-4.0" path="*." verb="*" type="System.Web.Handlers.TransferRequestHandler" preCondition="integratedMode,runtimeVersionv4.0" />
    </handlers>
        <httpProtocol>
          <customHeaders>
            <add name="Access-Control-Allow-Origin" value="*" />
          </customHeaders>
        </httpProtocol>
  </system.webServer>
```

	<runtime> <assemblyBinding xmlns="urn:schemas-microsoft-com:asm.v1"> <dependentAssembly> <assemblyIdentity name="System.Web.Helpers" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-2.0.0.0" newVersion="2.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="System.Web.Mvc" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-4.0.0.0" newVersion="4.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="System.Web.WebPages" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-2.0.0.0" newVersion="2.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="WebGrease" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="0.0.0.0-1.3.0.0" newVersion="1.3.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="Newtonsoft.Json" publicKeyToken="30ad4fe6b2a6aeed" culture="neutral" /> <bindingRedirect oldVersion="0.0.0.0-6.0.0.0" newVersion="6.0.0.0" /> </dependentAssembly> </assemblyBinding> </runtime> </configuration> <!--ProjectGuid: D84968F2-96A7-4E7D-B2BC-62BE07CF7DC1-->
Zawartość pliku konfiguracyjnego #2	IIS - C:\RepoAV\Services\RepositoryAccess\Web.config <?xml version="1.0" encoding="utf-8"?> <!-- For more information on how to configure your ASP.NET application, please visit http://go.microsoft.com/fwlink/?LinkId=169433 --> <configuration> <configSections>

```
    <section name="loggingConfiguration"
type="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.Logg
ingSettings, Microsoft.Practices.EnterpriseLibrary.Logging" />
    <section name="repositoryNodesConfiguration"
type="PSNC.RepoAV.Services.RepositoryAccess.RepositoryNodesConfi
gurationSection, RepositoryAccess" />
  </configSections>
  <repositoryNodesConfiguration>
    <redirection enabled="true" smartRedirectForRequestHeader="EP-
CDN-REQ" />
    <repositoryNodes thisNodeId="4">
      <add id="4" enabled="true" address="10.90.54.14" />
      <add id="1" enabled="true" address="10.90.54.11" />
      <add id="2" enabled="true" address="10.90.54.12" />
      <add id="3" enabled="true" address="10.90.54.13" />
            <add id="5" enabled="true" address="10.90.54.15" />
            <add id="6" enabled="true" address="10.90.54.16" />
    </repositoryNodes>
  </repositoryNodesConfiguration>
  <appSettings>
    <add key="NetworkAddressAndMaskWithoutAuthorization"
value="10.90.54.0/24" />
    <add key="ResponseCachingTime" value="5" />
  </appSettings>
  <connectionStrings>
    <add name="ContentDB" connectionString="Data
Source=(local);Initial Catalog=MaterialFormatDB;Integrated
Security=SSPI" providerName="System.Data.SqlClient" />
    <add name="RepoDB" connectionString="Data Source=ORAV1;Initial
Catalog=RepDB;User ID=sa;Password=ore!infomex!2017"
providerName="System.Data.SqlClient" />
  </connectionStrings>
  <loggingConfiguration name="Logging Application Block"
tracingEnabled="false" defaultCategory=""
logWarningsWhenNoCategoriesMatch="false">
    <listeners>
      <add fileName="C:\RepoAV\Logs\RepositoryAccess2.log" header=""
footer="" formatter="Text Formatter"
listenerDataType="Microsoft.Practices.EnterpriseLibrary.Logging.Config
uration.FlatFileTraceListenerData,
Microsoft.Practices.EnterpriseLibrary.Logging"
traceOutputOptions="None"
type="Microsoft.Practices.EnterpriseLibrary.Logging.TraceListeners.Flat
FileTraceListener, Microsoft.Practices.EnterpriseLibrary.Logging"
name="FlatFile TraceListener" />
    </listeners>
    <formatters>
```

```
      <add template="[{timestamp(local:yyyy-MM-dd
HH:mm:ss.ffff)}]&#xA;[{severity}][{category}] {message}&#xA;"
type="Microsoft.Practices.EnterpriseLibrary.Logging.Formatters.TextFor
matter, Microsoft.Practices.EnterpriseLibrary.Logging" name="Text
Formatter" />
    </formatters>
    <categorySources>
     <add switchValue="Information" name="General">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </add>
     <add switchValue="Information" name="RepositoryAccess">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </add>
    </categorySources>
    <specialSources>
     <allEvents switchValue="All" name="All Events" />
     <notProcessed switchValue="All" name="Unprocessed Category" />
     <errors switchValue="All" name="Logging Errors &amp;
Warnings">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </errors>
    </specialSources>
  </loggingConfiguration>
  <system.web>
    <compilation debug="true" targetFramework="4.5" />
    <httpRuntime targetFramework="4.5" />
  </system.web>
  <system.webServer>
    <modules>
     <remove name="DownloadModule" />
     <add name="DownloadModule"
type="PSNC.RepoAV.Services.RepositoryAccess.DownloadModule" />
    </modules>
      <httpProtocol>
        <customHeaders>
          <add name="Access-Control-Allow-Headers" value="X-
Requested-With" />
        </customHeaders>
      </httpProtocol>
  </system.webServer>
</configuration>
<!--ProjectGuid: 1D806E58-A600-49DC-9D31-E49BE52FB371-->
```

I. Nazwa serwera: orav5

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	orav5
IP / Domena	10.90.54.15
OS	Windows Server 2012 R2 Standard
RAM:	32 GB
vCPU	10
Przestrzeń	450 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Otwarte Repozytorium A/V

2. Konfiguracja .hosts

10.90.54.31	rt.epo.pl
10.90.54.250	av.epo.pl
10.90.54.250	av.epodreczniki.pl

3. Interfejsy sieciowe

LAN 10.90.54.15

4. Zainstalowane usługi

a. Bazy danych

Parametr	Wartość
Lokalizacja instalacji	MSSQL Server 2014 Express MaterialFormatDB C:\Program Files\Microsoft SQL Server\MSSQL12.MSSQLSERVER\MSSQL\DATA

b. RepoAV

Parametr	Wartość
Lokalizacja instalacji	C:\RepoAV E:\Storage
Zawartość pliku konfiguracyjnego #1	IIS - C:\RepoAV\Services\RepAPI\Web.config <?xml version="1.0" encoding="utf-8"?> <!-- For more information on how to configure your ASP.NET application, please visit http://go.microsoft.com/fwlink/?LinkId=169433 --> <configuration> <configSections> <section name="loggingConfiguration" type="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.LoggingSettings, Microsoft.Practices.EnterpriseLibrary.Logging" /> </configSections> <connectionStrings> <add name="DefaultConnection" providerName="System.Data.SqlClient" connectionString="Data Source=ORAV1;Initial Catalog=RepDB;User ID=sa;Password=ore!infomex!2017" /> </connectionStrings> <appSettings> <add key="webpages:Version" value="2.0.0.0" /> <add key="webpages:Enabled" value="false" /> <add key="PreserveLoginUrl" value="true" /> <add key="ClientValidationEnabled" value="true" /> <add key="UnobtrusiveJavaScriptEnabled" value="true" /> </appSettings> <loggingConfiguration name="Logging Application Block" tracingEnabled="false" defaultCategory="" logWarningsWhenNoCategoriesMatch="false"> <listeners> <add fileName="C:\RepoAV\Logs\RepAPI.log" header="" footer="" formatter="Text Formatter" listenerDataType="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.FlatFileTraceListenerData, Microsoft.Practices.EnterpriseLibrary.Logging" traceOutputOptions="None" type="Microsoft.Practices.EnterpriseLibrary.Logging.TraceListeners.FlatFileTraceListener, Microsoft.Practices.EnterpriseLibrary.Logging" name="FlatFile TraceListener" /> </listeners> <formatters> <add template="[{timestamp(local:yyyy-MM-dd HH:mm:ss.ffff)}]
[{severity}][{category}] {message}
"

```
type="Microsoft.Practices.EnterpriseLibrary.Logging.Formatters.TextFormatter, Microsoft.Practices.EnterpriseLibrary.Logging" name="Text Formatter" />
    </formatters>
    <categorySources>
      <add switchValue="Information" name="General">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </add>
      <add switchValue="Information" name="RepositoryAccess">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </add>
    </categorySources>
    <specialSources>
      <allEvents switchValue="All" name="All Events" />
      <notProcessed switchValue="All" name="Unprocessed Category" />
      <errors switchValue="All" name="Logging Errors &amp;amp; Warnings">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </errors>
    </specialSources>
  </loggingConfiguration>
  <system.web>
    <compilation debug="true" targetFramework="4.5" />
    <httpRuntime targetFramework="4.5" />
    <authentication mode="None" />
  </system.web>
  <system.webServer>
    <validation validateIntegratedModeConfiguration="false" />
    <handlers>
      <remove name="ExtensionlessUrlHandler-Integrated-4.0" />
      <remove name="OPTIONSVerbHandler" />
      <remove name="TRACEVerbHandler" />
      <add name="ExtensionlessUrlHandler-Integrated-4.0" path="*." verb="*" type="System.Web.Handlers.TransferRequestHandler" preCondition="integratedMode,runtimeVersionv4.0" />
    </handlers>
        <httpProtocol>
          <customHeaders>
            <add name="Access-Control-Allow-Origin" value="*" />
          </customHeaders>
        </httpProtocol>
  </system.webServer>
```

	<runtime> <assemblyBinding xmlns="urn:schemas-microsoft-com:asm.v1"> <dependentAssembly> <assemblyIdentity name="System.Web.Helpers" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-2.0.0.0" newVersion="2.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="System.Web.Mvc" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-4.0.0.0" newVersion="4.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="System.Web.WebPages" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-2.0.0.0" newVersion="2.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="WebGrease" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="0.0.0.0-1.3.0.0" newVersion="1.3.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="Newtonsoft.Json" publicKeyToken="30ad4fe6b2a6aeed" culture="neutral" /> <bindingRedirect oldVersion="0.0.0.0-6.0.0.0" newVersion="6.0.0.0" /> </dependentAssembly> </assemblyBinding> </runtime> </configuration> <!--ProjectGuid: D84968F2-96A7-4E7D-B2BC-62BE07CF7DC1-->
Zawartość pliku konfiguracyjnego #2	IIS - C:\RepoAV\Services\RepositoryAccess\Web.config <?xml version="1.0" encoding="utf-8"?> <!-- For more information on how to configure your ASP.NET application, please visit http://go.microsoft.com/fwlink/?LinkId=169433 --> <configuration> <configSections>

```
    <section name="loggingConfiguration"
type="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.Logg
ingSettings, Microsoft.Practices.EnterpriseLibrary.Logging" />
    <section name="repositoryNodesConfiguration"
type="PSNC.RepoAV.Services.RepositoryAccess.RepositoryNodesConfi
gurationSection, RepositoryAccess" />
  </configSections>
  <repositoryNodesConfiguration>
    <redirection enabled="true" smartRedirectForRequestHeader="EP-
CDN-REQ" />
    <repositoryNodes thisNodeId="5">
      <add id="5" enabled="true" address="10.90.54.15" />
      <add id="1" enabled="true" address="10.90.54.11" />
      <add id="2" enabled="true" address="10.90.54.12" />
      <add id="3" enabled="true" address="10.90.54.13" />
      <add id="4" enabled="true" address="10.90.54.14" />
            <add id="6" enabled="true" address="10.90.54.16" />
    </repositoryNodes>
  </repositoryNodesConfiguration>
  <appSettings>
    <add key="NetworkAddressAndMaskWithoutAuthorization"
value="10.90.54.0/24" />
    <add key="ResponseCachingTime" value="5" />
  </appSettings>
  <connectionStrings>
    <add name="ContentDB" connectionString="Data
Source=(local);Initial Catalog=MaterialFormatDB;Integrated
Security=SSPI" providerName="System.Data.SqlClient" />
    <add name="RepoDB" connectionString="Data Source=ORAV1;Initial
Catalog=RepDB;User ID=sa;Password=ore!infomex!2017"
providerName="System.Data.SqlClient" />
  </connectionStrings>
  <loggingConfiguration name="Logging Application Block"
tracingEnabled="false" defaultCategory=""
logWarningsWhenNoCategoriesMatch="false">
    <listeners>
      <add fileName="C:\RepoAV\Logs\RepositoryAccess2.log" header=""
footer="" formatter="Text Formatter"
listenerDataType="Microsoft.Practices.EnterpriseLibrary.Logging.Config
uration.FlatFileTraceListenerData,
Microsoft.Practices.EnterpriseLibrary.Logging"
traceOutputOptions="None"
type="Microsoft.Practices.EnterpriseLibrary.Logging.TraceListeners.Flat
FileTraceListener, Microsoft.Practices.EnterpriseLibrary.Logging"
name="FlatFile TraceListener" />
    </listeners>
    <formatters>
```

```
      <add template="[{timestamp(local:yyyy-MM-dd
HH:mm:ss.ffff)}]&#xA;[{severity}][{category}] {message}&#xA;"
type="Microsoft.Practices.EnterpriseLibrary.Logging.Formatters.TextFor
matter, Microsoft.Practices.EnterpriseLibrary.Logging" name="Text
Formatter" />
    </formatters>
    <categorySources>
     <add switchValue="Information" name="General">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </add>
     <add switchValue="Information" name="RepositoryAccess">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </add>
    </categorySources>
    <specialSources>
     <allEvents switchValue="All" name="All Events" />
     <notProcessed switchValue="All" name="Unprocessed Category" />
     <errors switchValue="All" name="Logging Errors &amp;
Warnings">
      <listeners>
       <add name="FlatFile TraceListener" />
      </listeners>
     </errors>
    </specialSources>
  </loggingConfiguration>
  <system.web>
    <compilation debug="true" targetFramework="4.5" />
    <httpRuntime targetFramework="4.5" />
  </system.web>
  <system.webServer>
    <modules>
     <remove name="DownloadModule" />
     <add name="DownloadModule"
type="PSNC.RepoAV.Services.RepositoryAccess.DownloadModule" />
    </modules>
      <httpProtocol>
        <customHeaders>
          <add name="Access-Control-Allow-Headers" value="X-
Requested-With" />
        </customHeaders>
      </httpProtocol>
  </system.webServer>
</configuration>
<!--ProjectGuid: 1D806E58-A600-49DC-9D31-E49BE52FB371-->
```

I. Nazwa serwera: orav6

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	orav6
IP / Domena	10.90.54.16
OS	Windows Server 2012 R2 Standard
RAM:	32 GB
vCPU	10
Przestrzeń	450 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Otwarte Repozytorium A/V

2. Konfiguracja .hosts

```
10.90.54.31     rt.epo.pl
10.90.54.250    av.epo.pl
10.90.54.250    av.epodreczniki.pl
```

3. Interfejsy sieciowe

LAN 10.90.54.16

4. Zainstalowane usługi

   a. Bazy danych

Parametr	Wartość
Lokalizacja instalacji	MSSQL Server 2014 Express MaterialFormatDB C:\Program Files\Microsoft SQL Server\MSSQL12.MSSQLSERVER\MSSQL\DATA

   b. RepoAV

Parametr	Wartość
Lokalizacja instalacji	C:\RepoAV E:\Storage
Zawartość pliku konfiguracyjnego #1	IIS - C:\RepoAV\Services\RepAPI\Web.config <?xml version="1.0" encoding="utf-8"?> <!-- For more information on how to configure your ASP.NET application, please visit http://go.microsoft.com/fwlink/?LinkId=169433 --> <configuration> <configSections> <section name="loggingConfiguration" type="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.LoggingSettings, Microsoft.Practices.EnterpriseLibrary.Logging" /> </configSections> <connectionStrings> <add name="DefaultConnection" providerName="System.Data.SqlClient" connectionString="Data Source=ORAV1;Initial Catalog=RepDB;User ID=sa;Password=ore!infomex!2017" /> </connectionStrings> <appSettings> <add key="webpages:Version" value="2.0.0.0" /> <add key="webpages:Enabled" value="false" /> <add key="PreserveLoginUrl" value="true" /> <add key="ClientValidationEnabled" value="true" /> <add key="UnobtrusiveJavaScriptEnabled" value="true" /> </appSettings> <loggingConfiguration name="Logging Application Block" tracingEnabled="false" defaultCategory="" logWarningsWhenNoCategoriesMatch="false"> <listeners> <add fileName="C:\RepoAV\Logs\RepAPI.log" header="" footer="" formatter="Text Formatter" listenerDataType="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.FlatFileTraceListenerData, Microsoft.Practices.EnterpriseLibrary.Logging" traceOutputOptions="None" type="Microsoft.Practices.EnterpriseLibrary.Logging.TraceListeners.FlatFileTraceListener, Microsoft.Practices.EnterpriseLibrary.Logging" name="FlatFile TraceListener" /> </listeners> <formatters> <add template="[{timestamp(local:yyyy-MM-dd HH:mm:ss.ffff)}]
[{severity}][{category}] {message}
" type="Microsoft.Practices.EnterpriseLibrary.Logging.Formatters.TextFor

```
matter, Microsoft.Practices.EnterpriseLibrary.Logging" name="Text
Formatter" />
    </formatters>
    <categorySources>
      <add switchValue="Information" name="General">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </add>
      <add switchValue="Information" name="RepositoryAccess">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </add>
    </categorySources>
    <specialSources>
      <allEvents switchValue="All" name="All Events" />
      <notProcessed switchValue="All" name="Unprocessed Category" />
      <errors switchValue="All" name="Logging Errors &amp;
Warnings">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </errors>
    </specialSources>
  </loggingConfiguration>
  <system.web>
    <compilation debug="true" targetFramework="4.5" />
    <httpRuntime targetFramework="4.5" />
    <authentication mode="None" />
  </system.web>
  <system.webServer>
    <validation validateIntegratedModeConfiguration="false" />
    <handlers>
      <remove name="ExtensionlessUrlHandler-Integrated-4.0" />
      <remove name="OPTIONSVerbHandler" />
      <remove name="TRACEVerbHandler" />
      <add name="ExtensionlessUrlHandler-Integrated-4.0" path="*."
verb="*" type="System.Web.Handlers.TransferRequestHandler"
preCondition="integratedMode,runtimeVersionv4.0" />
    </handlers>
        <httpProtocol>
          <customHeaders>
            <add name="Access-Control-Allow-Origin" value="*" />
          </customHeaders>
        </httpProtocol>
  </system.webServer>
  <runtime>
```

	<assemblyBinding xmlns="urn:schemas-microsoft-com:asm.v1"> <dependentAssembly> <assemblyIdentity name="System.Web.Helpers" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-2.0.0.0" newVersion="2.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="System.Web.Mvc" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-4.0.0.0" newVersion="4.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="System.Web.WebPages" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="1.0.0.0-2.0.0.0" newVersion="2.0.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="WebGrease" publicKeyToken="31bf3856ad364e35" /> <bindingRedirect oldVersion="0.0.0.0-1.3.0.0" newVersion="1.3.0.0" /> </dependentAssembly> <dependentAssembly> <assemblyIdentity name="Newtonsoft.Json" publicKeyToken="30ad4fe6b2a6aeed" culture="neutral" /> <bindingRedirect oldVersion="0.0.0.0-6.0.0.0" newVersion="6.0.0.0" /> </dependentAssembly> </assemblyBinding> </runtime> </configuration> <!--ProjectGuid: D84968F2-96A7-4E7D-B2BC-62BE07CF7DC1-->
Zawartość pliku konfiguracyjnego #2	IIS - C:\RepoAV\Services\RepositoryAccess\Web.config <?xml version="1.0" encoding="utf-8"?> <!-- For more information on how to configure your ASP.NET application, please visit http://go.microsoft.com/fwlink/?LinkId=169433 --> <configuration> <configSections> <section name="loggingConfiguration" type="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.LoggingSettings, Microsoft.Practices.EnterpriseLibrary.Logging" />

```
    <section name="repositoryNodesConfiguration"
type="PSNC.RepoAV.Services.RepositoryAccess.RepositoryNodesConfi
gurationSection, RepositoryAccess" />
  </configSections>
  <repositoryNodesConfiguration>
    <redirection enabled="true" smartRedirectForRequestHeader="EP-
CDN-REQ" />
    <repositoryNodes thisNodeId="6">
      <add id="6" enabled="true" address="10.90.54.16" />
      <add id="1" enabled="true" address="10.90.54.11" />
      <add id="2" enabled="true" address="10.90.54.12" />
      <add id="3" enabled="true" address="10.90.54.13" />
      <add id="4" enabled="true" address="10.90.54.14" />
      <add id="5" enabled="true" address="10.90.54.15" />
    </repositoryNodes>
  </repositoryNodesConfiguration>
  <appSettings>
    <add key="NetworkAddressAndMaskWithoutAuthorization"
value="10.90.54.0/24" />
    <add key="ResponseCachingTime" value="5" />
  </appSettings>
  <connectionStrings>
    <add name="ContentDB" connectionString="Data
Source=(local);Initial Catalog=MaterialFormatDB;Integrated
Security=SSPI" providerName="System.Data.SqlClient" />
    <add name="RepoDB" connectionString="Data Source=ORAV1;Initial
Catalog=RepDB;User ID=sa;Password=ore!infomex!2017"
providerName="System.Data.SqlClient" />
  </connectionStrings>
  <loggingConfiguration name="Logging Application Block"
tracingEnabled="false" defaultCategory=""
logWarningsWhenNoCategoriesMatch="false">
    <listeners>
      <add fileName="C:\RepoAV\Logs\RepositoryAccess2.log" header=""
footer="" formatter="Text Formatter"
listenerDataType="Microsoft.Practices.EnterpriseLibrary.Logging.Config
uration.FlatFileTraceListenerData,
Microsoft.Practices.EnterpriseLibrary.Logging"
traceOutputOptions="None"
type="Microsoft.Practices.EnterpriseLibrary.Logging.TraceListeners.Flat
FileTraceListener, Microsoft.Practices.EnterpriseLibrary.Logging"
name="FlatFile TraceListener" />
    </listeners>
    <formatters>
      <add template="[{timestamp(local:yyyy-MM-dd
HH:mm:ss.ffff)}]&#xA;[{severity}][{category}] {message}&#xA;"
type="Microsoft.Practices.EnterpriseLibrary.Logging.Formatters.TextFor
```

```
matter, Microsoft.Practices.EnterpriseLibrary.Logging" name="Text
Formatter" />
    </formatters>
    <categorySources>
      <add switchValue="Information" name="General">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </add>
      <add switchValue="Information" name="RepositoryAccess">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </add>
    </categorySources>
    <specialSources>
      <allEvents switchValue="All" name="All Events" />
      <notProcessed switchValue="All" name="Unprocessed Category" />
      <errors switchValue="All" name="Logging Errors &amp;
Warnings">
        <listeners>
          <add name="FlatFile TraceListener" />
        </listeners>
      </errors>
    </specialSources>
  </loggingConfiguration>
  <system.web>
    <compilation debug="true" targetFramework="4.5" />
    <httpRuntime targetFramework="4.5" />
  </system.web>
  <system.webServer>
    <modules>
      <remove name="DownloadModule" />
      <add name="DownloadModule"
type="PSNC.RepoAV.Services.RepositoryAccess.DownloadModule" />
    </modules>
        <httpProtocol>
          <customHeaders>
            <add name="Access-Control-Allow-Headers" value="X-
Requested-With" />
          </customHeaders>
        </httpProtocol>
  </system.webServer>
</configuration>
<!--ProjectGuid: 1D806E58-A600-49DC-9D31-E49BE52FB371-->
```

I. Nazwa serwera: ort1

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	ort1
IP / Domena	10.90.54.31
OS	CentOS 7
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	2300 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Aplikacyjny, bazodanowy

2. Konfiguracja .hosts

```
127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4 ort1
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6 ort1
10.90.54.31 ort1 rt.epodreczniki.pl
10.90.54.31 rt.epo.pl rt.epodreczniki.pl epodreczniki.pcss.pl
10.90.54.91 rabbit.epo.pl

10.90.54.200 epo.pl www.epo.pl epodreczniki.pl www.epodreczniki.pl
10.90.54.200 api.epo.pl www.api.epo.pl www.api.epodreczniki.pl
10.90.54.200 dev.epo.pl www.dev.epo.pl www.dev.epodreczniki.pl
10.90.54.200 static.epo.pl www.static.epo.pl static.epodreczniki.pl www.static.epodreczniki.pl
10.90.54.200 content.epo.pl www.content.epo.pl content.epodreczniki.pl www.content.epodreczniki.pl
10.90.54.200 preview.epo.pl www.preview.epo.pl preview.epodreczniki.pl www.preview.epodreczniki.pl
10.90.54.200 edit.epo.pl www.edit.epo.pl edit.epodreczniki.pl www.edit.epodreczniki.pl
10.90.54.200 repo.epo.pl www.repo.epo.pl repo.epodreczniki.pl www.repo.epodreczniki.pl
10.90.54.200 user.epo.pl www.user.epo.pl user.epodreczniki.pl www.user.epodreczniki.pl
10.90.54.200 search.epo.pl www.search.epo.pl www.search.epodreczniki.pl www.search.epodreczniki.pl
10.90.54.250 av.epo.pl www.av.epo.pl av.epodreczniki.pl www.av.epodreczniki.pl
```

3. Interfejsy sieciowe

eth0 10.90.54.31

4. Zainstalowane usługi

a. TomCat

Parametr	Wartość
Lokalizacja instalacji	/etc/tomcat
Zawartość pliku konfiguracyjnego #1	/etc/tomcat/tomcat.conf # System-wide configuration file for tomcat services # This will be loaded by systemd as an environment file, # so please keep the syntax. For shell expansion support # place your custom files as /etc/tomcat/conf.d/*.conf # # There are 2 "classes" of startup behavior in this package. # The old one, the default service named tomcat.service. # The new named instances are called tomcat@instance.service. # # Use this file to change default values for all services. # Change the service specific ones to affect only one service. # For tomcat.service it's /etc/sysconfig/tomcat, for # tomcat@instance it's /etc/sysconfig/tomcat@instance.

	# This variable is used to figure out if config is loaded or not.
	TOMCAT_CFG_LOADED="1"
	
	# In new-style instances, if CATALINA_BASE isn't specified, it will
	# be constructed by joining TOMCATS_BASE and NAME.
	TOMCATS_BASE="/var/lib/tomcats/"
	
	# Where your java installation lives
	JAVA_HOME="/usr/lib/jvm/jre"
	
	# Where your tomcat installation lives
	CATALINA_HOME="/usr/share/tomcat"
	
	# System-wide tmp
	CATALINA_TMPDIR="/var/cache/tomcat/temp"
	
	# You can pass some parameters to java here if you wish to
	#JAVA_OPTS="-Xminf0.1 -Xmaxf0.3"
	
	# Use JAVA_OPTS to set java.library.path for libtcnative.so
	#JAVA_OPTS="-Djava.library.path=/usr/lib"
	
	# You can change your tomcat locale here
	#LANG="en_US"
	
	# Run tomcat under the Java Security Manager
	SECURITY_MANAGER="false"
	
	# Time to wait in seconds, before killing process
	# TODO(stingray): does nothing, fix.
	# SHUTDOWN_WAIT="30"
	
	# If you wish to further customize your tomcat environment,
	# put your own definitions here
	# (i.e. LD_LIBRARY_PATH for some jdbc drivers)

Zawartość pliku konfiguracyjnego #2	/etc/tomcat/server.xml <?xml version='1.0' encoding='utf-8'?> <!-- Licensed to the Apache Software Foundation (ASF) under one or more contributor license agreements. See the NOTICE file distributed with this work for additional information regarding copyright ownership. The ASF licenses this file to You under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at http://www.apache.org/licenses/LICENSE-2.0 Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License. --> <!-- Note: A "Server" is not itself a "Container", so you may not define subcomponents such as "Valves" at this level. Documentation at /docs/config/server.html --> <Server port="8005" shutdown="SHUTDOWN"> <Listener className="org.apache.catalina.startup.VersionLoggerListener" /> <!-- Security listener. Documentation at /docs/config/listeners.html <Listener className="org.apache.catalina.security.SecurityListener" /> --> <!--APR library loader. Documentation at /docs/apr.html --> <Listener className="org.apache.catalina.core.AprLifecycleListener" SSLEngine="on" /> <!--Initialize Jasper prior to webapps are loaded. Documentation at /docs/jasper-howto.html --> <Listener className="org.apache.catalina.core.JasperListener" /> <!-- Prevent memory leaks due to use of particular java/javax APIs--> <Listener className="org.apache.catalina.core.JreMemoryLeakPreventionListener" /> <Listener className="org.apache.catalina.mbeans.GlobalResourcesLifecycleListener" /> <Listener className="org.apache.catalina.core.ThreadLocalLeakPreventionListener" />

```
<!-- Global JNDI resources
     Documentation at /docs/jndi-resources-howto.html
-->
<GlobalNamingResources>
  <!-- Editable user database that can also be used by
       UserDatabaseRealm to authenticate users
  -->
  <Resource name="UserDatabase" auth="Container"
            type="org.apache.catalina.UserDatabase"
            description="User database that can be updated and saved"
            factory="org.apache.catalina.users.MemoryUserDatabaseFactory"
            pathname="conf/tomcat-users.xml" />
</GlobalNamingResources>

<!-- A "Service" is a collection of one or more "Connectors" that share
     a single "Container" Note:  A "Service" is not itself a "Container",
     so you may not define subcomponents such as "Valves" at this level.
     Documentation at /docs/config/service.html
 -->
<Service name="Catalina">

  <!--The connectors can use a shared executor, you can define one or more named thread pools-->
  <!--
  <Executor name="tomcatThreadPool" namePrefix="catalina-exec-"
    maxThreads="150" minSpareThreads="4"/>
  -->


  <!-- A "Connector" represents an endpoint by which requests are received
       and responses are returned. Documentation at :
       Java HTTP Connector: /docs/config/http.html (blocking & non-blocking)
       Java AJP  Connector: /docs/config/ajp.html
       APR (HTTP/AJP) Connector: /docs/apr.html
       Define a non-SSL HTTP/1.1 Connector on port 8080
  -->
  <Connector port="8080" protocol="HTTP/1.1"
             connectionTimeout="20000"
             redirectPort="8443" />
  <!-- A "Connector" using the shared thread pool-->
  <!--
  <Connector executor="tomcatThreadPool"
             port="8080" protocol="HTTP/1.1"
             connectionTimeout="20000"
             redirectPort="8443" />
  -->
  <!-- Define a SSL HTTP/1.1 Connector on port 8443
       This connector uses the BIO implementation that requires the JSSE
       style configuration. When using the APR/native implementation, the
       OpenSSL style configuration is required as described in the APR/native
       documentation -->
  <!--
  <Connector port="8443" protocol="org.apache.coyote.http11.Http11Protocol"
             maxThreads="150" SSLEnabled="true" scheme="https" secure="true"
             clientAuth="false" sslProtocol="TLS" />
  -->

  <!-- Define an AJP 1.3 Connector on port 8009 -->
  <Connector port="8009" protocol="AJP/1.3" redirectPort="8443" />


  <!-- An Engine represents the entry point (within Catalina) that processes
       every request.  The Engine implementation for Tomcat stand alone
       analyzes the HTTP headers included with the request, and passes them
       on to the appropriate Host (virtual host).
       Documentation at /docs/config/engine.html -->

  <!-- You should set jvmRoute to support load-balancing via AJP ie :
  <Engine name="Catalina" defaultHost="localhost" jvmRoute="jvm1">
  -->
  <Engine name="Catalina" defaultHost="localhost">

    <!--For clustering, please take a look at documentation at:
        /docs/cluster-howto.html  (simple how to)
```

	/docs/config/cluster.html (reference documentation) --> <!-- <Cluster className="org.apache.catalina.ha.tcp.SimpleTcpCluster"/> --> <!-- Use the LockOutRealm to prevent attempts to guess user passwords via a brute-force attack --> <Realm className="org.apache.catalina.realm.LockOutRealm"> <!-- This Realm uses the UserDatabase configured in the global JNDI resources under the key "UserDatabase". Any edits that are performed against this UserDatabase are immediately available for use by the Realm. --> <Realm className="org.apache.catalina.realm.UserDatabaseRealm" resourceName="UserDatabase"/> </Realm> <Host name="localhost" appBase="webapps" unpackWARs="true" autoDeploy="true"> <!-- SingleSignOn valve, share authentication between web applications Documentation at: /docs/config/valve.html --> <!-- <Valve className="org.apache.catalina.authenticator.SingleSignOn" /> --> <!-- Access log processes all example. Documentation at: /docs/config/valve.html Note: The pattern used is equivalent to using pattern="common" --> <Valve className="org.apache.catalina.valves.AccessLogValve" directory="logs" prefix="localhost_access_log." suffix=".txt" pattern="%h %l %u %t "%r" %s %b" /> </Host> </Engine> </Service> </Server>

b. Dlibra

Parametr	Wartość
Lokalizacja instalacji	/root/dlibra-server
Zawartość pliku konfiguracyjnego #1	/root/dlibra-server/conf/server.xml <?xml version="1.0"?> <!DOCTYPE dlibra [<------><!ELEMENT dlibra (systemServicesUrl?,databaseConfigurationFile?,smtpConfigurationFile?,luceneConfigurationFile?,extensions Directory,userAuthenticatorClass,service-list)> <------><!ATTLIST dlibra. <------><------>serverHost CDATA #REQUIRED. <------><------>serverPort CDATA #REQUIRED> <------><!ELEMENT systemServicesUrl (#PCDATA)><><------> <------><!ELEMENT databaseConfigurationFile (#PCDATA)> <------><!ELEMENT smtpConfigurationFile (#PCDATA)><----> <------><!ELEMENT luceneConfigurationFile (#PCDATA)> <------><!ELEMENT extensionsDirectory (#PCDATA)> <------><!ELEMENT userAuthenticatorClass (#PCDATA)> <------><!ELEMENT service EMPTY> <------><!ATTLIST service. <------><------>serviceDescription CDATA #IMPLIED<-----><------> <------><------>serviceType NMTOKEN #IMPLIED <------><------>serviceVersion CDATA #IMPLIED <------><------>serviceFactoryClass NMTOKEN #IMPLIED.<-><------> <------><------>password CDATA #IMPLIED> <------><!ELEMENT service-list (service+)> <------>]> <dlibra serverHost="rt.epo.pl" serverPort="10051"> <------><!-- The entry below is useful only when dLibra server is splitted into <------><------>multiple machines. If it is commented out then value.

```
<------><------>//127.0.0.1:10051/ss.
<------><------>is used instead..
<------>-->
<------><!--.
<------><------><systemServicesUrl>
<------><------>//system-services-service-hostname-here:port-of-the-service-here/ss
<------><------></systemServicesUrl>.
<------>-->
<------><databaseConfigurationFile>
<------><------>database.properties
<------></databaseConfigurationFile>
<------><!-- The entry below is useful only when server requires e-mail sending..
<------><------>It is not needed by default.
<------>-->
<------><!--.
<------><------><smtpConfigurationFile>smtp.properties</smtpConfigurationFile>
<------>-->
<------><luceneConfigurationFile>lucene.properties</luceneConfigurationFile>
<------><extensionsDirectory>extensions</extensionsDirectory>
<------><userAuthenticatorClass>pl.psnc.dlibra.service.DlibraUserAuthenticator</userAuthenticatorClass>
<------><service-list>
<------><------><service serviceDescription="System services" serviceType="ss"
<------><------><------>serviceVersion="1"
<------><------><------>serviceFactoryClass="pl.psnc.dlibra.system.server.SystemServicesFactory"
<------><------><------>password="aaa" />
<------><------><service serviceDescription="User interface" serviceType="ui"
<------><------><------>serviceVersion="1"
<------><------><------>serviceFactoryClass="pl.psnc.dlibra.system.server.UserInterfaceFactory"
<------><------><------>password="aaa" />
<------><------><service serviceDescription="User server" serviceType="us"
<------><------><------>serviceVersion="1"
<------><------><------>serviceFactoryClass="pl.psnc.dlibra.user.server.UserServerFactory"
<------><------><------>password="aaa" />
<------><------><service serviceDescription="Event manager" serviceType="em"
<------><------><------>serviceVersion="1"
<------><------><------>serviceFactoryClass="pl.psnc.dlibra.event.server.EventServerFactory"
<------><------><------>password="aaa" />
<------><------><service serviceDescription="Content server" serviceType="cs"
<------><------><------>serviceVersion="1"
<------><------><------>serviceFactoryClass="pl.psnc.dlibra.content.server.ContentServerFactory"
<------><------><------>password="aaa" />
<------><------><service serviceDescription="Metadata server" serviceType="ms"
<------><------><------>serviceVersion="1"
<------><------><------>serviceFactoryClass="pl.psnc.dlibra.metadata.server.MetadataServerFactory"
<------><------><------>password="aaa" />
<------><------><service serviceDescription="Search server" serviceType="se"
<------><------><------>serviceVersion="1"
<------><------><------>serviceFactoryClass="pl.psnc.dlibra.search.server.SearchServerFactory"
<------><------><------>password="aaa" />
<------><------><service serviceDescription="Index server" serviceType="is"
<------><------><------>serviceVersion="1"
<------><------><------>serviceFactoryClass="pl.psnc.dlibra.index.server.IndexServerFactory"
<------><------><------>password="aaa" />
<------><------><service serviceDescription="Profile Provider" serviceType="pp"
<------><------><------>serviceVersion="1"
<------><------><------>serviceFactoryClass="pl.psnc.dlibra.profile.server.ProfileProviderFactory"
<------><------><------>password="aaa" />
<------><------><service serviceDescription="Message server" serviceType="me"
<------><------><------>serviceVersion="1"
<------><------><------>serviceFactoryClass="pl.psnc.dlibra.message.server.MessageServerFactory"
<------><------><------>password="aaa" />
<------><------><!-- IMPORTANT: This service should be in each services cluster. -->
<------><------><service serviceDescription="dLibra JMX Management Service"
<------><------><------>serviceType="mx" serviceVersion="1"
<------><------><------>serviceFactoryClass="pl.psnc.dlibra.server.jmx.DLibraJMXServiceFactory"
<------><------><------>password="aaa" />
<------><------><service serviceDescription="ep"
<------><------><------>serviceType="ep"
<------><------><------>serviceVersion="1"
<------><------><------>serviceFactoryClass="pl.psnc.ep.rt.server.EPServiceFactory"
<------><------><------>password="aaa" />
<------></service-list>
</dlibra>
```

I. Nazwa serwera: ore-rfv

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	ore-rfv
IP / Domena	91.224.214.98, 10.90.54.254, 172.16.55.1
OS	Ubuntu 16.04
RAM:	6 GB
vCPU	4 vCPU
Przestrzeń	50
Zewnętrzne IP:	Tak 91.224.214.98
Funkcja serwera:	router, firewall, vpn, balancer, ips

2. Konfiguracja .hosts

```
127.0.0.1 localhost
91.224.214.98 ore-rfv

10.90.54.250 av.epo.pl
10.90.54.201 app1
10.90.54.202 app2
10.90.54.203 app3

# The following lines are desirable for IPv6 capable hosts
::1 ip6-localhost ip6-loopback
fe00::0 ip6-localnet
ff00::0 ip6-mcastprefix
ff02::1 ip6-allnodes
ff02::2 ip6-allrouters
ff02::3 ip6-allhosts
```

3. Interfejsy sieciowe

ens3 91.224.214.98
ens4 10.90.54.254
tun0 172.16.55.1

4. Zainstalowane usługi
    a. HaProxy

Parametr	Wartość
Lokalizacja instalacji	/etc/haproxy/haproxy.cfg
Zawartość pliku konfiguracyjnego #1	global log /dev/log local0 chroot /var/lib/haproxy stats socket /run/haproxy/admin.sock mode 660 level admin stats timeout 30s daemon # transparent proxy wymaga roota # user haproxy # group haproxy # SSL CIPHERS ca-base /etc/ssl/certs crt-base /etc/ssl/private ssl-default-bind-ciphers ECDH+AESGCM:DH+AESGCM:ECDH+AES256:DH+AES256:ECDH+AES128:DH+AES:ECDH+3DES:DH+3DES:RSA+AESGCM:RSA+AES:RSA+3DES:!aNULL:!MD5:!DSS ssl-default-bind-options no-sslv3 defaults log global mode http option httplog option dontlognull timeout connect 5000 timeout client 50000 timeout server 50000 errorfile 400 /etc/haproxy/errors/400.http errorfile 403 /etc/haproxy/errors/403.http errorfile 408 /etc/haproxy/errors/408.http errorfile 500 /etc/haproxy/errors/500.http errorfile 502 /etc/haproxy/errors/502.http errorfile 503 /etc/haproxy/errors/503.http errorfile 504 /etc/haproxy/errors/504.http ## ## statystyki ##

```
listen ore-stats.epodreczniki.pl
  bind 91.224.214.98:8000 transparent
  mode http
  stats uri /haproxy?stats

##
## epodreczniki.pl - glowna strona portalu
##

frontend front-epodreczniki.pl
   bind 91.224.214.99:80 transparent
   bind 91.224.214.99:443 transparent ssl crt
/etc/haproxy/certs/wildcard.epodreczniki.pl.pem
   mode http

# CORS dla portalu, dodanie nagłówków doodpowiedzi i do requestów.
   http-response add-header Access-Control-Allow-Origin
%[capture.req.hdr(0)] if { capture.req.hdr(0) -m found }
   rspadd Access-Control-Allow-Headers:\ Origin,\ Accept,\ X-Requested-
With if { capture.req.hdr(0) -m found }

 ## redirect auth to ssl
   acl      is_user.epodreczniki.pl     hdr(host)  -i user.epodreczniki.pl
   acl      is_auth_login               path_beg   -i /auth/login
   redirect scheme  https if is_user.epodreczniki.pl is_auth_login !{ ssl_fc }

   ## redirect all to https
   #redirect scheme  https  if !{ ssl_fc }

   ## default backend
   default_backend back-epodreczniki.pl

backend back-epodreczniki.pl
   source 0.0.0.0 usesrc clientip
   option httpclose
   option forwardfor
   balance roundrobin

   server app1 10.90.54.201:80 check fall 3 rise 2
   server app2 10.90.54.202:80 check fall 3 rise 2
   server app3 10.90.54.203:80 check fall 3 rise 2

# Czwórka, wyłączona z klastra.

   #server app4 10.90.54.204:80 check
```

```
##
## epodreczniki.pl - internal proxy for portal
##
listen front-int-epodreczniki.pl
   bind     10.90.54.200:80
   balance  roundrobin
   mode     http
   option   forwardfor

   server app1 10.90.54.201:80 check fall 3 rise 2
   server app2 10.90.54.202:80 check fall 3 rise 2
   server app3 10.90.54.203:80 check fall 3 rise 2

# Czwórka wyłączona z klastra na prośbę Michała Sołtyska - dla niego tym
momencie to pre-prod.

   #server app4 10.90.54.204:80 check

##
## av.epodreczniki.pl do serwerow windows po http
##

frontend front-av.epodreczniki.pl
   bind 91.224.214.100:80 transparent
   bind 91.224.214.100:443 transparent ssl crt
/etc/haproxy/certs/wildcard.epodreczniki.pl.pem
   mode http
   option forwardfor
   capture request header origin len 203
   http-request set-header EP-CDN-REQ AV

   # CORS

# redirect all to http
   #redirect scheme  https  if !{ ssl_fc }

   acl repo_api url_beg /RepApi
   #acl ore_allowed_net src 10.90.54.0/24

   acl host_av1 hdr(host) -i av1.epodreczniki.pl
   acl host_av2 hdr(host) -i av2.epodreczniki.pl
   acl host_av3 hdr(host) -i av3.epodreczniki.pl
   acl host_av4 hdr(host) -i av4.epodreczniki.pl
   acl host_av5 hdr(host) -i av5.epodreczniki.pl
   acl host_av6 hdr(host) -i av6.epodreczniki.pl

   acl host_av1_get url_end &sn=1
   acl host_av2_get url_end &sn=2
```

```
    acl host_av3_get url_end &sn=3
    acl host_av4_get url_end &sn=4
    acl host_av5_get url_end &sn=5
    acl host_av6_get url_end &sn=6

    #block if repo_api !ore_allowed_net

    use_backend back-av1.epodreczniki.pl if host_av1 or host_av1_get
    use_backend back-av2.epodreczniki.pl if host_av2 or host_av2_get
    use_backend back-av3.epodreczniki.pl if host_av3 or host_av3_get
    use_backend back-av4.epodreczniki.pl if host_av4 or host_av4_get
    use_backend back-av5.epodreczniki.pl if host_av5 or host_av5_get
    use_backend back-av6.epodreczniki.pl if host_av6 or host_av6_get

    default_backend back-av.epodreczniki.pl


backend back-av.epodreczniki.pl
    source 0.0.0.0 usesrc clientip
    option httpclose
    option forwardfor
    balance source
    server av.01 10.90.54.11:80 check
    server av.02 10.90.54.12:80 check
    server av.03 10.90.54.13:80 check
    server av.04 10.90.54.14:80 check
    server av.05 10.90.54.15:80 check
    server av.06 10.90.54.16:80 check

backend back-av1.epodreczniki.pl
    source 0.0.0.0 usesrc clientip
    option httpclose
    option forwardfor
    balance roundrobin
    server av.01 10.90.54.11:80 check

backend back-av2.epodreczniki.pl
    source 0.0.0.0 usesrc clientip
    option httpclose
    option forwardfor
    balance roundrobin
    server av.02 10.90.54.12:80 check

backend back-av3.epodreczniki.pl
    source 0.0.0.0 usesrc clientip
    option httpclose
    option forwardfor
    balance roundrobin
```

```
    server av.03 10.90.54.13:80 check

backend back-av4.epodreczniki.pl
    source 0.0.0.0 usesrc clientip
    option httpclose
    option forwardfor
    balance roundrobin
    server av.04 10.90.54.14:80 check

backend back-av5.epodreczniki.pl
    source 0.0.0.0 usesrc clientip
    option httpclose
    option forwardfor
    balance roundrobin
    server av.04 10.90.54.15:80 check

backend back-av6.epodreczniki.pl
    source 0.0.0.0 usesrc clientip
    option httpclose
    option forwardfor
    balance roundrobin
    server av.04 10.90.54.16:80 check

##
## rt.epodreczniki.pl po SSL i HTTP
##
frontend rt.epodreczniki.pl-ssl
    bind 91.224.214.101:443 transparent
    option tcplog
    mode tcp
    default_backend rt-ssl

backend rt-ssl
    source 0.0.0.0 usesrc clientip
    mode tcp
    balance roundrobin
    option ssl-hello-chk
    server rt.31 10.90.54.31:443 check


frontend rt.epodreczniki.pl-http
    bind 91.224.214.101:80 transparent
    mode http
    default_backend rt-http

backend rt-http
    source 0.0.0.0 usesrc clientip
    option httpclose
```

	option forwardfor balance roundrobin server rt.31 10.90.54.31:80 check ## ## av.epo.pl lokalnie #### lokalny, w razie w trzeba skopiowac ustawienia z wyzej, aby obslugiwal przekierowanie url_end $sn=1-6 jak w wanowym frontcie. ## frontend front-av.epo.pl bind 10.90.54.250:80 #transparent mode http stats enable stats uri /haproxy?stats default_backend back-av.epo.pl backend back-av.epo.pl # source 0.0.0.0 usesrc clientip option httpclose option forwardfor balance roundrobin server av.epo.pl.01 10.90.54.11:80 check server av.epo.pl.02 10.90.54.12:80 check server av.epo.pl.03 10.90.54.13:80 check server av.epo.pl.04 10.90.54.14:80 check server av.epo.pl.05 10.90.54.15:80 check server av.epo.pl.06 10.90.54.16:80 check #listen manager.av.epo.pl # bind 10.90.54.250:8088 transparent # mode http # balance roundrobin # option httpclose # option forwardfor # server av.epo.pl.01 10.90.54.11:8088 check # server av.epo.pl.02 10.90.54.12:8088 check backup # server av.epo.pl.03 10.90.54.13:8088 check backup
Zawartość pliku konfiguracyjnego #2	<wklej treść pliku>

b. OpenVPN

Parametr	Wartość
Lokalizacja instalacji	/etc/openvpn/openvpn.conf
Zawartość pliku konfiguracyjnego #1	;local a.b.c.d port 1195 proto udp dev tun ca /etc/openvpn/easyrsa/keys/ca.crt cert /etc/openvpn/easyrsa/keys/ore-ovpn.crt key /etc/openvpn/easyrsa/keys/ore-ovpn.key # This file should be kept secret dh /etc/openvpn/easyrsa/keys/dh2048.pem #crl-verify /etc/openvpn/easyrsa/keys/crl.pem server 172.16.55.0 255.255.255.0 ifconfig-pool-persist ipp.txt push "route 10.90.54.0 255.255.255.0" ;push "dhcp-option DNS 208.67.222.222" ;push "dhcp-option DNS 208.67.220.220" ;client-to-client ;duplicate-cn keepalive 10 120 ;tls-auth ta.key 0 # This file is secret cipher AES-128-CBC # AES comp-lzo ;max-clients 100 ;user nobody ;group nogroup persist-key persist-tun status /var/log/openvpn/openvpn-status.log

	;log openvpn.log log-append /var/log/openvpn/openvpn.log verb 3 ;mute 20

I. Nazwa serwera: skt1

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	skt1
IP / Domena	10.90.54.51
OS	CentOS 7
RAM:	32 GB
vCPU	10 v CPU
Przestrzeń	1300 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Kodowanie treści

2. Konfiguracja .hosts

```
127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6

10.90.54.31 ort1 rt.epodreczniki.pl
10.90.54.91 rabbit.epo.pl

10.90.54.200 epo.pl www.epo.pl epodreczniki.pl www.epodreczniki.pl
10.90.54.200 api.epo.pl www.api.epo.pl api.epodreczniki.pl www.api.epodreczniki.pl
10.90.54.200 static.epo.pl www.static.epo.pl static.epodreczniki.pl www.static.epodreczniki.pl
10.90.54.200 content.epo.pl www.content.epo.pl content.epodreczniki.pl www.content.epodreczniki.pl
10.90.54.200 preview.epo.pl www.preview.epo.pl preview.epodreczniki.pl www.preview.epodreczniki.pl
10.90.54.200 edit.epo.pl www.edit.epo.pl edit.epodreczniki.pl www.edit.epodreczniki.pl
10.90.54.200 repo.epo.pl www.repo.epo.pl repo.epodreczniki.pl www.repo.epodreczniki.pl
10.90.54.200 user.epo.pl www.user.epo.pl user.epodreczniki.pl www.user.epodreczniki.pl
10.90.54.200 search.epo.pl www.search.epo.pl www.search.epodreczniki.pl www.search.epodreczniki.pl
10.90.54.250 av.epo.pl www.av.epo.pl av.epodreczniki.pl www.av.epodreczniki.pl
10.90.54.31  rt.epo.pl rt.epodreczniki.pl epodreczniki.pcss.pl
```

3. Interfejsy sieciowe

eth0 10.90.54.51

4. Zainstalowane usługi

a. Apache-maven

Parametr	Wartość
Lokalizacja instalacji	/opt/apache-maven/

Zawartość pliku konfiguracyjnego #1

```
<?xml version="1.0" encoding="UTF-8"?>


<!--
 | This is the configuration file for Maven. It can be specified at two levels:
 |
 | 1. User Level. This settings.xml file provides configuration for a single user,
 |                and is normally provided in ${user.home}/.m2/settings.xml.
 |
 |                NOTE: This location can be overridden with the CLI option:
 |
 |                -s /path/to/user/settings.xml
 |
 | 2. Global Level. This settings.xml file provides configuration for all Maven
 |                users on a machine (assuming they're all using the same Maven
 |                installation). It's normally provided in
 |                ${maven.home}/conf/settings.xml.
 |
 |                NOTE: This location can be overridden with the CLI option:
 |
 |                -gs /path/to/global/settings.xml
 |
 | The sections in this sample file are intended to give you a running start at
 | getting the most out of your Maven installation. Where appropriate, the default
 | values (values used when the setting is not specified) are provided.
 |
 |-->
<settings xmlns="http://maven.apache.org/SETTINGS/1.0.0"
          xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
          xsi:schemaLocation="http://maven.apache.org/SETTINGS/1.0.0
http://maven.apache.org/xsd/settings-1.0.0.xsd">
  <!-- localRepository
   | The path to the local repository maven will use to store artifacts.
   |
   | Default: ${user.home}/.m2/repository
  <localRepository>/path/to/local/repo</localRepository>
  -->

  <!-- interactiveMode
   | This will determine whether maven prompts you when it needs input. If set to false,
   | maven will use a sensible default value, perhaps based on some other setting, for
   | the parameter in question.
   |
   | Default: true
  <interactiveMode>true</interactiveMode>
  -->

  <!-- offline
   | Determines whether maven should attempt to connect to the network when executing a build.
   | This will have an effect on artifact downloads, artifact deployment, and others.
   |
   | Default: false
  <offline>false</offline>
  -->

  <!-- pluginGroups
   | This is a list of additional group identifiers that will be searched when resolving plugins by their
prefix, i.e.
   | when invoking a command line like "mvn prefix:goal". Maven will automatically add the group
identifiers
   | "org.apache.maven.plugins" and "org.codehaus.mojo" if these are not already contained in the list.
   |-->
  <pluginGroups>
    <!-- pluginGroup
     | Specifies a further group identifier to use for plugin lookup.
    <pluginGroup>com.your.plugins</pluginGroup>
    -->
  </pluginGroups>

  <!-- proxies
   | This is a list of proxies which can be used on this machine to connect to the network.
   | Unless otherwise specified (by system property or command-line switch), the first proxy
   | specification in this list marked as active will be used.
```

```
  |-->
 <proxies>
  <!-- proxy
   | Specification for one proxy, to be used in connecting to the network.
   |
  <proxy>
   <id>optional</id>
   <active>true</active>
   <protocol>http</protocol>
   <username>proxyuser</username>
   <password>proxypass</password>
   <host>proxy.host.net</host>
   <port>80</port>
   <nonProxyHosts>local.net|some.host.com</nonProxyHosts>
  </proxy>
  -->
 </proxies>

 <!-- servers
  | This is a list of authentication profiles, keyed by the server-id used within the system.
  | Authentication profiles can be used whenever maven must make a connection to a remote server.
  |-->
 <servers>
  <!-- server
   | Specifies the authentication information to use when connecting to a particular server, identified
by
   | a unique name within the system (referred to by the 'id' attribute below).
   |
   | NOTE: You should either specify username/password OR privateKey/passphrase, since these
pairings are
   |       used together.
   |
  <server>
   <id>deploymentRepo</id>
   <username>repouser</username>
   <password>repopwd</password>
  </server>
  -->

  <!-- Another sample, using keys to authenticate.
  <server>
   <id>siteServer</id>
   <privateKey>/path/to/private/key</privateKey>
   <passphrase>optional; leave empty if not used.</passphrase>
  </server>
  -->
 </servers>

 <!-- mirrors
  | This is a list of mirrors to be used in downloading artifacts from remote repositories.
  |
  | It works like this: a POM may declare a repository to use in resolving certain artifacts.
  | However, this repository may have problems with heavy traffic at times, so people have mirrored
  | it to several places.
  |
  | That repository definition will have a unique id, so we can create a mirror reference for that
  | repository, to be used as an alternate download site. The mirror site will be the preferred
  | server for that repository.
  |-->
 <mirrors>
  <!-- mirror
   | Specifies a repository mirror site to use instead of a given repository. The repository that
   | this mirror serves has an ID that matches the mirrorOf element of this mirror. IDs are used
   | for inheritance and direct lookup purposes, and must be unique across the set of mirrors.
   |
  <mirror>
   <id>mirrorId</id>
   <mirrorOf>repositoryId</mirrorOf>
   <name>Human Readable Name for this Mirror.</name>
   <url>http://my.repository.com/repo/path</url>
  </mirror>
   -->
 </mirrors>
```

```
  <!-- profiles
   | This is a list of profiles which can be activated in a variety of ways, and which can modify
   | the build process. Profiles provided in the settings.xml are intended to provide local machine-
   | specific paths and repository locations which allow the build to work in the local environment.
   |
   | For example, if you have an integration testing plugin - like cactus - that needs to know where
   | your Tomcat instance is installed, you can provide a variable here such that the variable is
   | dereferenced during the build process to configure the cactus plugin.
   |
   | As noted above, profiles can be activated in a variety of ways. One way - the activeProfiles
   | section of this document (settings.xml) - will be discussed later. Another way essentially
   | relies on the detection of a system property, either matching a particular value for the property,
   | or merely testing its existence. Profiles can also be activated by JDK version prefix, where a
   | value of '1.4' might activate a profile when the build is executed on a JDK version of '1.4.2_07'.
   | Finally, the list of active profiles can be specified directly from the command line.
   |
   | NOTE: For profiles defined in the settings.xml, you are restricted to specifying only artifact
   |       repositories, plugin repositories, and free-form properties to be used as configuration
   |       variables for plugins in the POM.
   |
   |-->
  <profiles>
    <!-- profile
     | Specifies a set of introductions to the build process, to be activated using one or more of the
     | mechanisms described above. For inheritance purposes, and to activate profiles via
<activatedProfiles/>
     | or the command line, profiles have to have an ID that is unique.
     |
     | An encouraged best practice for profile identification is to use a consistent naming convention
     | for profiles, such as 'env-dev', 'env-test', 'env-production', 'user-jdcasey', 'user-brett', etc.
     | This will make it more intuitive to understand what the set of introduced profiles is attempting
     | to accomplish, particularly when you only have a list of profile id's for debug.
     |
     | This profile example uses the JDK version to trigger activation, and provides a JDK-specific repo.
    <profile>
      <id>jdk-1.4</id>

      <activation>
        <jdk>1.4</jdk>
      </activation>

      <repositories>
        <repository>
          <id>jdk14</id>
          <name>Repository for JDK 1.4 builds</name>
          <url>http://www.myhost.com/maven/jdk14</url>
          <layout>default</layout>
          <snapshotPolicy>always</snapshotPolicy>
        </repository>
      </repositories>
    </profile>
    -->

    <!--
     | Here is another profile, activated by the system property 'target-env' with a value of 'dev',
     | which provides a specific path to the Tomcat instance. To use this, your plugin configuration
     | might hypothetically look like:
     |
     | ...
     | <plugin>
     |   <groupId>org.myco.myplugins</groupId>
     |   <artifactId>myplugin</artifactId>
     |
     |   <configuration>
     |     <tomcatLocation>${tomcatPath}</tomcatLocation>
     |   </configuration>
     | </plugin>
     | ...
     |
     | NOTE: If you just wanted to inject this configuration whenever someone set 'target-env' to
     |       anything, you could just leave off the <value/> inside the activation-property.
     |
```

	<profile> <id>env-dev</id> <activation> <property> <name>target-env</name> <value>dev</value> </property> </activation> <properties> <tomcatPath>/path/to/tomcat/instance</tomcatPath> </properties> </profile> --> </profiles> <!-- activeProfiles | List of profiles that are active for all builds. | <activeProfiles> <activeProfile>alwaysActiveProfile</activeProfile> <activeProfile>anotherAlwaysActiveProfile</activeProfile> </activeProfiles> --> </settings>
Zawartość pliku konfiguracyjnego #2	org.slf4j.simpleLogger.defaultLogLevel=info org.slf4j.simpleLogger.showDateTime=false org.slf4j.simpleLogger.showThreadName=false org.slf4j.simpleLogger.showLogName=false org.slf4j.simpleLogger.logFile=System.out org.slf4j.simpleLogger.levelInBrackets=true org.slf4j.simpleLogger.log.Sisu=info org.slf4j.simpleLogger.warnLevelString=WARNING

I. Nazwa serwera: storage1

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	storage1
IP / Domena	10.90.54.211
OS	Ubuntu 16.04
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	650 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	plikowy

2. Konfiguracja .hosts

```
127.0.0.1 localhost storage1


# The following lines are desirable for IPv6 capable hosts
::1 ip6-localhost ip6-loopback
fe00::0 ip6-localnet
ff00::0 ip6-mcastprefix
ff02::1 ip6-allnodes
ff02::2 ip6-allrouters
ff02::3 ip6-allhosts


10.90.54.211 storage1
10.90.54.212 storage2
```

3. Interfejsy sieciowe

ens3 10.90.54.211

4. Zainstalowane usługi

a. Gluster

Parametr	Wartość
Lokalizacja instalacji	/etc/glusterfs
Zawartość pliku konfiguracyjnego #1	/var/lib/glusterd/vols/storage/storage.storage1.storage-brick0-data.vol volume storage-posix type storage/posix option volume-id ca65f018-4782-499c-aec5-e2c10c0fc007 option directory /storage/brick0/data end-volume volume storage-trash type features/trash option trash-internal-op off option brick-path /storage/brick0/data option trash-dir .trashcan subvolumes storage-posix end-volume

```
volume storage-changetimerecorder
    type features/changetimerecorder
    option record-counters off
    option ctr-enabled off
    option record-entry on
    option ctr_inode_heal_expire_period 300
    option ctr_hardlink_heal_expire_period 300
    option ctr_link_consistency off
    option record-exit off
    option db-path /storage/brick0/data/.glusterfs/
    option db-name data.db
    option hot-brick off
    option db-type sqlite3
    subvolumes storage-trash
end-volume

volume storage-changelog
    type features/changelog
    option changelog-barrier-timeout 120
    option changelog-dir /storage/brick0/data/.glusterfs/changelogs
    option changelog-brick /storage/brick0/data
    subvolumes storage-changetimerecorder
end-volume

volume storage-bitrot-stub
    type features/bitrot-stub
    option export /storage/brick0/data
    subvolumes storage-changelog
end-volume

volume storage-access-control
    type features/access-control
    subvolumes storage-bitrot-stub
end-volume

volume storage-locks
    type features/locks
    subvolumes storage-access-control
end-volume

volume storage-upcall
    type features/upcall
    option cache-invalidation off
    subvolumes storage-locks
end-volume

volume storage-io-threads
    type performance/io-threads
    subvolumes storage-upcall
end-volume

volume storage-marker
    type features/marker
    option inode-quota off
    option quota off
    option gsync-force-xtime off
    option xtime off
    option quota-version 0
    option timestamp-file /var/lib/glusterd/vols/storage/marker.tstamp
    option volume-uuid ca65f018-4782-499c-aec5-e2c10c0fc007
    subvolumes storage-io-threads
end-volume

volume storage-barrier
    type features/barrier
    option barrier-timeout 120
    option barrier disable
    subvolumes storage-marker
end-volume

volume storage-index
    type features/index
    option index-base /storage/brick0/data/.glusterfs/indices
```

	subvolumes storage-barrier end-volume volume storage-quota type features/quota option deem-statfs off option timeout 0 option server-quota off option volume-uuid storage subvolumes storage-index end-volume volume storage-worm type features/worm option worm off subvolumes storage-quota end-volume volume storage-read-only type features/read-only option read-only off subvolumes storage-worm end-volume volume /storage/brick0/data type debug/io-stats option count-fop-hits off option latency-measurement off subvolumes storage-read-only end-volume volume storage-server type protocol/server option auth.addr./storage/brick0/data.allow 10.90.54.* option auth.login.f9bb4230-6ff4-4edb-9803-b4602e405f52.password 4b6b1345-c8d8-473a-af9b-314b2da80545 option auth.login./storage/brick0/data.allow f9bb4230-6ff4-4edb-9803-b4602e405f52 option transport-type tcp subvolumes /storage/brick0/data end-volume
Zawartość pliku konfiguracyjnego #2	/var/lib/glusterd/vols/storage/storage.storage2.storage-brick0-data.vol volume storage-posix type storage/posix option volume-id ca65f018-4782-499c-aec5-e2c10c0fc007 option directory /storage/brick0/data end-volume volume storage-trash type features/trash option trash-internal-op off option brick-path /storage/brick0/data option trash-dir .trashcan subvolumes storage-posix end-volume volume storage-changetimerecorder type features/changetimerecorder option record-counters off option ctr-enabled off option record-entry on option ctr_inode_heal_expire_period 300 option ctr_hardlink_heal_expire_period 300 option ctr_link_consistency off option record-exit off option db-path /storage/brick0/data/.glusterfs/ option db-name data.db option hot-brick off option db-type sqlite3 subvolumes storage-trash

```
end-volume

volume storage-changelog
    type features/changelog
    option changelog-barrier-timeout 120
    option changelog-dir /storage/brick0/data/.glusterfs/changelogs
    option changelog-brick /storage/brick0/data
    subvolumes storage-changetimerecorder
end-volume

volume storage-bitrot-stub
    type features/bitrot-stub
    option export /storage/brick0/data
    subvolumes storage-changelog
end-volume

volume storage-access-control
    type features/access-control
    subvolumes storage-bitrot-stub
end-volume

volume storage-locks
    type features/locks
    subvolumes storage-access-control
end-volume

volume storage-upcall
    type features/upcall
    option cache-invalidation off
    subvolumes storage-locks
end-volume

volume storage-io-threads
    type performance/io-threads
    subvolumes storage-upcall
end-volume

volume storage-marker
    type features/marker
    option inode-quota off
    option quota off
    option gsync-force-xtime off
    option xtime off
    option quota-version 0
    option timestamp-file /var/lib/glusterd/vols/storage/marker.tstamp
    option volume-uuid ca65f018-4782-499c-aec5-e2c10c0fc007
    subvolumes storage-io-threads
end-volume

volume storage-barrier
    type features/barrier
    option barrier-timeout 120
    option barrier disable
    subvolumes storage-marker
end-volume

volume storage-index
    type features/index
    option index-base /storage/brick0/data/.glusterfs/indices
    subvolumes storage-barrier
end-volume

volume storage-quota
    type features/quota
    option deem-statfs off
    option timeout 0
    option server-quota off
    option volume-uuid storage
    subvolumes storage-index
end-volume

volume storage-worm
    type features/worm
```

	option worm off subvolumes storage-quota end-volume volume storage-read-only type features/read-only option read-only off subvolumes storage-worm end-volume volume /storage/brick0/data type debug/io-stats option count-fop-hits off option latency-measurement off subvolumes storage-read-only end-volume volume storage-server type protocol/server option auth.addr./storage/brick0/data.allow 10.90.54.* option auth.login.f9bb4230-6ff4-4edb-9803-b4602e405f52.password 4b6b1345-c8d8-473a-af9b-314b2da80545 option auth.login./storage/brick0/data.allow f9bb4230-6ff4-4edb-9803-b4602e405f52 option transport-type tcp subvolumes /storage/brick0/data end-volume
Zawartość pliku konfiguracyjnego #3	/var/lib/glusterd/vols/storage/storage.tcp-fuse.vol volume storage-client-0 type protocol/client option send-gids true option transport-type tcp option remote-subvolume /storage/brick0/data option remote-host storage1 option ping-timeout 42 end-volume volume storage-client-1 type protocol/client option send-gids true option transport-type tcp option remote-subvolume /storage/brick0/data option remote-host storage2 option ping-timeout 42 end-volume volume storage-replicate-0 type cluster/replicate subvolumes storage-client-0 storage-client-1 end-volume volume storage-dht type cluster/distribute subvolumes storage-replicate-0 end-volume volume storage-write-behind type performance/write-behind subvolumes storage-dht end-volume volume storage-read-ahead type performance/read-ahead subvolumes storage-write-behind end-volume volume storage-readdir-ahead type performance/readdir-ahead subvolumes storage-read-ahead end-volume volume storage-io-cache

	type performance/io-cache subvolumes storage-readdir-ahead end-volume volume storage-quick-read type performance/quick-read subvolumes storage-io-cache end-volume volume storage-open-behind type performance/open-behind subvolumes storage-quick-read end-volume volume storage-md-cache type performance/md-cache subvolumes storage-open-behind end-volume volume storage type debug/io-stats option count-fop-hits off option latency-measurement off subvolumes storage-md-cache end-volume
Zawartość pliku konfiguracyjnego #4	/var/lib/glusterd/vols/storage/trusted-storage.tcp-fuse.vol volume storage-client-0 type protocol/client option send-gids true option password 4b6b1345-c8d8-473a-af9b-314b2da80545 option username f9bb4230-6ff4-4edb-9803-b4602e405f52 option transport-type tcp option remote-subvolume /storage/brick0/data option remote-host storage1 option ping-timeout 42 end-volume volume storage-client-1 type protocol/client option send-gids true option password 4b6b1345-c8d8-473a-af9b-314b2da80545 option username f9bb4230-6ff4-4edb-9803-b4602e405f52 option transport-type tcp option remote-subvolume /storage/brick0/data option remote-host storage2 option ping-timeout 42 end-volume volume storage-replicate-0 type cluster/replicate subvolumes storage-client-0 storage-client-1 end-volume volume storage-dht type cluster/distribute subvolumes storage-replicate-0 end-volume volume storage-write-behind type performance/write-behind subvolumes storage-dht end-volume volume storage-read-ahead type performance/read-ahead subvolumes storage-write-behind end-volume volume storage-readdir-ahead type performance/readdir-ahead subvolumes storage-read-ahead

	end-volume volume storage-io-cache type performance/io-cache subvolumes storage-readdir-ahead end-volume volume storage-quick-read type performance/quick-read subvolumes storage-io-cache end-volume volume storage-open-behind type performance/open-behind subvolumes storage-quick-read end-volume volume storage-md-cache type performance/md-cache subvolumes storage-open-behind end-volume volume storage type debug/io-stats option count-fop-hits off option latency-measurement off subvolumes storage-md-cache end-volume
Zawartość pliku konfiguracyjnego #5	/etc/glusterfs/gluster-rsyslog-7.2.conf ##### gluster.conf ##### # ## If you want to log every message to the log file instead of ## intelligently suppressing repeated messages, set off to ## RepeatedMsgReduction. This change requires rsyslog restart ## (eg. run 'service rsyslog restart') # #$RepeatedMsgReduction off $RepeatedMsgReduction on $ModLoad mmjsonparse *.* :mmjsonparse: # ## The mmcount module provides the capability to count log messages by. ## severity or json property of given app-name. The count value is added ## into the log message as json property named 'mmcount' ## ## More info at http://www.rsyslog.com/doc/mmcount.html # #module(load="mmcount") #action(type="mmcount" appname="glusterd" key="!gf_code") # count each value of gf_code of appname glusterd #action(type="mmcount" appname="glusterfsd" key="!gf_code") # count each value of gf_code of appname glusterfsd #action(type="mmcount" appname="glusterfs" key="!gf_code") # count each value of gf_code of appname glusterfs template (name="Glusterfsd_dynLogFile" type="string" string="/var/log/glusterfs/bricks/%app-name%.log") template (name="Gluster_dynLogFile" type="string" string="/var/log/glusterfs/%app-name%.log") template(name="GLFS_template" type="list") { property(name="$!mmcount") constant(value="/") property(name="syslogfacility-text" caseConversion="upper") constant(value="/") property(name="syslogseverity-text" caseConversion="upper") constant(value=" ") constant(value="[") property(name="timereported" dateFormat="rfc3339") constant(value="] ")

	constant(value="[")
	property(name="$!gf_code")
	constant(value="] ")
	constant(value="[")
	property(name="$!gf_message")
	constant(value="] ")
	property(name="$!msg")
	constant(value="\n")
	}
	
	if $app-name contains 'glusterfsd' then {
	action(type="omfile"
	DynaFile="Glusterfsd_dynLogFile"
	Template="GLFS_template")
	stop
	}
	
	if $app-name contains 'gluster' then {
	action(type="omfile"
	DynaFile="Gluster_dynLogFile"
	Template="GLFS_template")
	stop
	}
	#
	## send email for every 50th mmcount
	#$ModLoad ommail
	#if $app-name == 'glusterfsd' and $!mmcount <> 0 and $!mmcount % 50 == 0 then {
	# $ActionMailSMTPServer smtp.example.com
	# $ActionMailFrom rsyslog@example.com
	# $ActionMailTo glusteradmin@example.com
	# $template mailSubject,"50th message of gf_code=9999 on %hostname%"
	# $template mailBody,"RSYSLOG Alert\r\nmsg='%msg%'"
	# $ActionMailSubject mailSubject
	# $ActionExecOnlyOnceEveryInterval 30
	# :ommail:;RSYSLOG_SyslogProtocol23Format
	#}
	#

Zawartość pliku konfiguracyjnego #6	/etc/glusterfs/glusterd.vol volume management type mgmt/glusterd option working-directory /var/lib/glusterd option transport-type socket,rdma option transport.socket.keepalive-time 10 option transport.socket.keepalive-interval 2 option transport.socket.read-fail-log off option ping-timeout 30 # option base-port 49152 end-volume

I. Nazwa serwera: storage2

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	storage2
IP / Domena	10.90.54.212
OS	Ubuntu 16.04
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	650 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	plikowy

2. Konfiguracja .hosts

```
127.0.0.1 localhost storage2

# The following lines are desirable for IPv6 capable hosts
::1 ip6-localhost ip6-loopback
fe00::0 ip6-localnet
ff00::0 ip6-mcastprefix
ff02::1 ip6-allnodes
ff02::2 ip6-allrouters
ff02::3 ip6-allhosts


10.90.54.211 storage1
10.90.54.212 storage2
```

3. Interfejsy sieciowe

ens3 10.90.54.212

4. Zainstalowane usługi

a. Gluster

Parametr	Wartość
Lokalizacja instalacji	/etc/glusterfs
Zawartość pliku konfiguracyjnego #1	/var/lib/glusterd/vols/storage/storage.storage1.storage-brick0-data.vol volume storage-posix type storage/posix option volume-id ca65f018-4782-499c-aec5-e2c10c0fc007 option directory /storage/brick0/data end-volume volume storage-trash type features/trash option trash-internal-op off option brick-path /storage/brick0/data option trash-dir .trashcan subvolumes storage-posix end-volume volume storage-changetimerecorder

```
    type features/changetimerecorder
    option record-counters off
    option ctr-enabled off
    option record-entry on
    option ctr_inode_heal_expire_period 300
    option ctr_hardlink_heal_expire_period 300
    option ctr_link_consistency off
    option record-exit off
    option db-path /storage/brick0/data/.glusterfs/
    option db-name data.db
    option hot-brick off
    option db-type sqlite3
    subvolumes storage-trash
end-volume

volume storage-changelog
    type features/changelog
    option changelog-barrier-timeout 120
    option changelog-dir /storage/brick0/data/.glusterfs/changelogs
    option changelog-brick /storage/brick0/data
    subvolumes storage-changetimerecorder
end-volume

volume storage-bitrot-stub
    type features/bitrot-stub
    option export /storage/brick0/data
    subvolumes storage-changelog
end-volume

volume storage-access-control
    type features/access-control
    subvolumes storage-bitrot-stub
end-volume

volume storage-locks
    type features/locks
    subvolumes storage-access-control
end-volume

volume storage-upcall
    type features/upcall
    option cache-invalidation off
    subvolumes storage-locks
end-volume

volume storage-io-threads
    type performance/io-threads
    subvolumes storage-upcall
end-volume

volume storage-marker
    type features/marker
    option inode-quota off
    option quota off
    option gsync-force-xtime off
    option xtime off
    option quota-version 0
    option timestamp-file /var/lib/glusterd/vols/storage/marker.tstamp
    option volume-uuid ca65f018-4782-499c-aec5-e2c10c0fc007
    subvolumes storage-io-threads
end-volume

volume storage-barrier
    type features/barrier
    option barrier-timeout 120
    option barrier disable
    subvolumes storage-marker
end-volume

volume storage-index
    type features/index
    option index-base /storage/brick0/data/.glusterfs/indices
    subvolumes storage-barrier
```

	end-volume volume storage-quota type features/quota option deem-statfs off option timeout 0 option server-quota off option volume-uuid storage subvolumes storage-index end-volume volume storage-worm type features/worm option worm off subvolumes storage-quota end-volume volume storage-read-only type features/read-only option read-only off subvolumes storage-worm end-volume volume /storage/brick0/data type debug/io-stats option count-fop-hits off option latency-measurement off subvolumes storage-read-only end-volume volume storage-server type protocol/server option auth.addr./storage/brick0/data.allow 10.90.54.* option auth.login.f9bb4230-6ff4-4edb-9803-b4602e405f52.password 4b6b1345-c8d8-473a-af9b-314b2da80545 option auth.login./storage/brick0/data.allow f9bb4230-6ff4-4edb-9803-b4602e405f52 option transport-type tcp subvolumes /storage/brick0/data end-volume
Zawartość pliku konfiguracyjnego #2	/var/lib/glusterd/vols/storage/storage.storage2.storage-brick0-data.vol volume storage-posix type storage/posix option volume-id ca65f018-4782-499c-aec5-e2c10c0fc007 option directory /storage/brick0/data end-volume volume storage-trash type features/trash option trash-internal-op off option brick-path /storage/brick0/data option trash-dir .trashcan subvolumes storage-posix end-volume volume storage-changetimerecorder type features/changetimerecorder option record-counters off option ctr-enabled off option record-entry on option ctr_inode_heal_expire_period 300 option ctr_hardlink_heal_expire_period 300 option ctr_link_consistency off option record-exit off option db-path /storage/brick0/data/.glusterfs/ option db-name data.db option hot-brick off option db-type sqlite3 subvolumes storage-trash end-volume

```
volume storage-changelog
    type features/changelog
    option changelog-barrier-timeout 120
    option changelog-dir /storage/brick0/data/.glusterfs/changelogs
    option changelog-brick /storage/brick0/data
    subvolumes storage-changetimerecorder
end-volume

volume storage-bitrot-stub
    type features/bitrot-stub
    option export /storage/brick0/data
    subvolumes storage-changelog
end-volume

volume storage-access-control
    type features/access-control
    subvolumes storage-bitrot-stub
end-volume

volume storage-locks
    type features/locks
    subvolumes storage-access-control
end-volume

volume storage-upcall
    type features/upcall
    option cache-invalidation off
    subvolumes storage-locks
end-volume

volume storage-io-threads
    type performance/io-threads
    subvolumes storage-upcall
end-volume

volume storage-marker
    type features/marker
    option inode-quota off
    option quota off
    option gsync-force-xtime off
    option xtime off
    option quota-version 0
    option timestamp-file /var/lib/glusterd/vols/storage/marker.tstamp
    option volume-uuid ca65f018-4782-499c-aec5-e2c10c0fc007
    subvolumes storage-io-threads
end-volume

volume storage-barrier
    type features/barrier
    option barrier-timeout 120
    option barrier disable
    subvolumes storage-marker
end-volume

volume storage-index
    type features/index
    option index-base /storage/brick0/data/.glusterfs/indices
    subvolumes storage-barrier
end-volume

volume storage-quota
    type features/quota
    option deem-statfs off
    option timeout 0
    option server-quota off
    option volume-uuid storage
    subvolumes storage-index
end-volume

volume storage-worm
    type features/worm
    option worm off
```

	subvolumes storage-quota end-volume volume storage-read-only type features/read-only option read-only off subvolumes storage-worm end-volume volume /storage/brick0/data type debug/io-stats option count-fop-hits off option latency-measurement off subvolumes storage-read-only end-volume volume storage-server type protocol/server option auth.addr./storage/brick0/data.allow 10.90.54.* option auth.login.f9bb4230-6ff4-4edb-9803-b4602e405f52.password 4b6b1345-c8d8-473a-af9b-314b2da80545 option auth.login./storage/brick0/data.allow f9bb4230-6ff4-4edb-9803-b4602e405f52 option transport-type tcp subvolumes /storage/brick0/data end-volume
Zawartość pliku konfiguracyjnego #3	/var/lib/glusterd/vols/storage/storage.tcp-fuse.vol volume storage-client-0 type protocol/client option send-gids true option transport-type tcp option remote-subvolume /storage/brick0/data option remote-host storage1 option ping-timeout 42 end-volume volume storage-client-1 type protocol/client option send-gids true option transport-type tcp option remote-subvolume /storage/brick0/data option remote-host storage2 option ping-timeout 42 end-volume volume storage-replicate-0 type cluster/replicate subvolumes storage-client-0 storage-client-1 end-volume volume storage-dht type cluster/distribute subvolumes storage-replicate-0 end-volume volume storage-write-behind type performance/write-behind subvolumes storage-dht end-volume volume storage-read-ahead type performance/read-ahead subvolumes storage-write-behind end-volume volume storage-readdir-ahead type performance/readdir-ahead subvolumes storage-read-ahead end-volume volume storage-io-cache type performance/io-cache

	subvolumes storage-readdir-ahead end-volume volume storage-quick-read type performance/quick-read subvolumes storage-io-cache end-volume volume storage-open-behind type performance/open-behind subvolumes storage-quick-read end-volume volume storage-md-cache type performance/md-cache subvolumes storage-open-behind end-volume volume storage type debug/io-stats option count-fop-hits off option latency-measurement off subvolumes storage-md-cache end-volume
Zawartość pliku konfiguracyjnego #4	/var/lib/glusterd/vols/storage/trusted-storage.tcp-fuse.vol volume storage-client-0 type protocol/client option send-gids true option password 4b6b1345-c8d8-473a-af9b-314b2da80545 option username f9bb4230-6ff4-4edb-9803-b4602e405f52 option transport-type tcp option remote-subvolume /storage/brick0/data option remote-host storage1 option ping-timeout 42 end-volume volume storage-client-1 type protocol/client option send-gids true option password 4b6b1345-c8d8-473a-af9b-314b2da80545 option username f9bb4230-6ff4-4edb-9803-b4602e405f52 option transport-type tcp option remote-subvolume /storage/brick0/data option remote-host storage2 option ping-timeout 42 end-volume volume storage-replicate-0 type cluster/replicate subvolumes storage-client-0 storage-client-1 end-volume volume storage-dht type cluster/distribute subvolumes storage-replicate-0 end-volume volume storage-write-behind type performance/write-behind subvolumes storage-dht end-volume volume storage-read-ahead type performance/read-ahead subvolumes storage-write-behind end-volume volume storage-readdir-ahead type performance/readdir-ahead subvolumes storage-read-ahead end-volume

	volume storage-io-cache type performance/io-cache subvolumes storage-readdir-ahead end-volume volume storage-quick-read type performance/quick-read subvolumes storage-io-cache end-volume volume storage-open-behind type performance/open-behind subvolumes storage-quick-read end-volume volume storage-md-cache type performance/md-cache subvolumes storage-open-behind end-volume volume storage type debug/io-stats option count-fop-hits off option latency-measurement off subvolumes storage-md-cache end-volume
Zawartość pliku konfiguracyjnego #5	/etc/glusterfs/gluster-rsyslog-7.2.conf ##### gluster.conf ##### # ## If you want to log every message to the log file instead of ## intelligently suppressing repeated messages, set off to ## RepeatedMsgReduction. This change requires rsyslog restart ## (eg. run 'service rsyslog restart') # #$RepeatedMsgReduction off $RepeatedMsgReduction on $ModLoad mmjsonparse *.* :mmjsonparse: # ## The mmcount module provides the capability to count log messages by. ## severity or json property of given app-name. The count value is added ## into the log message as json property named 'mmcount' ## ## More info at http://www.rsyslog.com/doc/mmcount.html # #module(load="mmcount") #action(type="mmcount" appname="glusterd" key="!gf_code") # count each value of gf_code of appname glusterd #action(type="mmcount" appname="glusterfsd" key="!gf_code") # count each value of gf_code of appname glusterfsd #action(type="mmcount" appname="glusterfs" key="!gf_code") # count each value of gf_code of appname glusterfs template (name="Glusterfsd_dynLogFile" type="string" string="/var/log/glusterfs/bricks/%app-name%.log") template (name="Gluster_dynLogFile" type="string" string="/var/log/glusterfs/%app-name%.log") template(name="GLFS_template" type="list") { property(name="$!mmcount") constant(value="/") property(name="syslogfacility-text" caseConversion="upper") constant(value="/") property(name="syslogseverity-text" caseConversion="upper") constant(value=" ") constant(value="[") property(name="timereported" dateFormat="rfc3339") constant(value="] ") constant(value="[")

	property(name="$!gf_code")
	constant(value="] ")
	constant(value="[")
	property(name="$!gf_message")
	constant(value="] ")
	property(name="$!msg")
	constant(value="\n")
	}
	
	if $app-name contains 'glusterfsd' then {
	action(type="omfile"
	DynaFile="Glusterfsd_dynLogFile"
	Template="GLFS_template")
	stop
	}
	
	if $app-name contains 'gluster' then {
	action(type="omfile"
	DynaFile="Gluster_dynLogFile"
	Template="GLFS_template")
	stop
	}
	#
	## send email for every 50th mmcount
	#$ModLoad ommail
	#if $app-name == 'glusterfsd' and $!mmcount <> 0 and $!mmcount % 50 == 0 then {
	# $ActionMailSMTPServer smtp.example.com
	# $ActionMailFrom rsyslog@example.com
	# $ActionMailTo glusteradmin@example.com
	# $template mailSubject,"50th message of gf_code=9999 on %hostname%"
	# $template mailBody,"RSYSLOG Alert\r\nmsg='%msg%'"
	# $ActionMailSubject mailSubject
	# $ActionExecOnlyOnceEveryInterval 30
	# :ommail:;RSYSLOG_SyslogProtocol23Format
	#}
	#

Zawartość pliku konfiguracyjnego #6	/etc/glusterfs/glusterd.vol volume management type mgmt/glusterd option working-directory /var/lib/glusterd option transport-type socket,rdma option transport.socket.keepalive-time 10 option transport.socket.keepalive-interval 2 option transport.socket.read-fail-log off option ping-timeout 30 # option base-port 49152 end-volume

I. Nazwa serwera: swkpk1

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	swkpk1
IP / Domena	10.90.54.91
OS	CentOS 7
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	200 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Wymiana komunikatów

2. Konfiguracja .hosts

```
127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4 ort1
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6 ort1

10.90.54.31 ort1 rt.epo.pl rt.epodreczniki.pl epodreczniki.pcss.pl
10.90.54.91 rabbit.epo.pl

10.90.54.200 epo.pl www.epo.pl epodreczniki.pl www.epodreczniki.pl
10.90.54.200 api.epo.pl www.api.epo.pl api.epodreczniki.pl www.api.epodreczniki.pl
10.90.54.200 static.epo.pl www.static.epo.pl static.epodreczniki.pl www.static.epodreczniki.pl
10.90.54.200 content.epo.pl www.content.epo.pl content.epodreczniki.pl www.content.epodreczniki.pl
10.90.54.200 preview.epo.pl www.preview.epo.pl preview.epodreczniki.pl www.preview.epodreczniki.pl
10.90.54.200 edit.epo.pl www.edit.epo.pl edit.epodreczniki.pl www.edit.epodreczniki.pl
10.90.54.200 repo.epo.pl www.repo.epo.pl repo.epodreczniki.pl www.repo.epodreczniki.pl
10.90.54.200 user.epo.pl www.user.epo.pl user.epodreczniki.pl www.user.epodreczniki.pl
10.90.54.200 search.epo.pl www.search.epo.pl www.search.epodreczniki.pl www.search.epodreczniki.pl
10.90.54.250 av.epo.pl www.av.epo.pl av.epodreczniki.pl www.av.epodreczniki.pl
```

3. Interfejsy sieciowe

eth0 10.90.54.91

4. Zainstalowane usługi

a. RabbitMQ

Parametr	Wartość
Lokalizacja instalacji	/etc/rabbitmq/
Zawartość pliku konfiguracyjnego #1	/etc/rabbitmq/rabitmq.config %% -*- mode: erlang -*- %% ---------- %% RabbitMQ Sample Configuration File. %% %% See http://www.rabbitmq.com/configure.html for details. %% ---------- [{rabbit, [%% %% Network Connectivity %% ==================== %% %% By default, RabbitMQ will listen on all interfaces, using

```
%% the standard (reserved) AMQP port.
%%
%% {tcp_listeners, [5672]},

%% To listen on a specific interface, provide a tuple of {IpAddress, Port}.
%% For example, to listen only on localhost for both IPv4 and IPv6:
%%
%% {tcp_listeners, [{"127.0.0.1", 5672},
%%                  {"::1",       5672}]},

%% SSL listeners are configured in the same fashion as TCP listeners,
%% including the option to control the choice of interface.
%%
%% {ssl_listeners, [5671]},

%% Log levels (currently just used for connection logging).
%% One of 'info', 'warning', 'error' or 'none', in decreasing order
%% of verbosity. Defaults to 'info'.
%%
{log_levels, [{connection, info}]},

%% Set to 'true' to perform reverse DNS lookups when accepting a
%% connection. Hostnames will then be shown instead of IP addresses
%% in rabbitmqctl and the management plugin.
%%
{reverse_dns_lookups, false}

%%
%% Security / AAA
%% ==============
%%

%% Configuring SSL.
%% See http://www.rabbitmq.com/ssl.html for full documentation.
%%
%% {ssl_options, [{cacertfile,           "/path/to/testca/cacert.pem"},
%%                {certfile,             "/path/to/server/cert.pem"},
%%                {keyfile,              "/path/to/server/key.pem"},
%%                {verify,               verify_peer},
%%                {fail_if_no_peer_cert, false}]},

%% Choose the available SASL mechanism(s) to expose.
%% The two default (built in) mechanisms are 'PLAIN' and
%% 'AMQPLAIN'. Additional mechanisms can be added via
%% plugins.
%%
%% See http://www.rabbitmq.com/authentication.html for more details.
%%
%% {auth_mechanisms, ['PLAIN', 'AMQPLAIN']},

%% Select an authentication database to use. RabbitMQ comes bundled
%% with a built-in auth-database, based on mnesia.
%%
%% {auth_backends, [rabbit_auth_backend_internal]},

%% Configurations supporting the rabbitmq_auth_mechanism_ssl and
%% rabbitmq_auth_backend_ldap plugins.
%%
%% NB: These options require that the relevant plugin is enabled.
%% See http://www.rabbitmq.com/plugins.html for further details.

%% The RabbitMQ-auth-mechanism-ssl plugin makes it possible to
%% authenticate a user based on the client's SSL certificate.
%%
%% To use auth-mechanism-ssl, add to or replace the auth_mechanisms
%% list with the entry 'EXTERNAL'.
%%
%% {auth_mechanisms, ['EXTERNAL']},

%% The rabbitmq_auth_backend_ldap plugin allows the broker to
%% perform authentication and authorisation by deferring to an
%% external LDAP server.
%%
```

```
%% For more information about configuring the LDAP backend, see
%% http://www.rabbitmq.com/ldap.html.
%%
%% Enable the LDAP auth backend by adding to or replacing the
%% auth_backends entry:
%%
%% {auth_backends, [rabbit_auth_backend_ldap]},

%% This pertains to both the rabbitmq_auth_mechanism_ssl plugin and
%% STOMP ssl_cert_login configurations. See the rabbitmq_stomp
%% configuration section later in this fail and the README in
%% https://github.com/rabbitmq/rabbitmq-auth-mechanism-ssl for further
%% details.
%%
%% To use the SSL cert's CN instead of its DN as the username
%%
%% {ssl_cert_login_from, common_name},

%%
%% Default User / VHost
%% ====================
%%

%% On first start RabbitMQ will create a vhost and a user. These
%% config items control what gets created. See
%% http://www.rabbitmq.com/access-control.html for further
%% information about vhosts and access control.
%%
%% {default_vhost,       <<"/">>},
%% {default_user,        <<"guest">>},
%% {default_pass,        <<"guest">>},
%% {default_permissions, [<<".*">>, <<".*">>, <<".*">>]},

%% Tags for default user
%%
%% For more details about tags, see the documentation for the
%% Management Plugin at http://www.rabbitmq.com/management.html.
%%
%% {default_user_tags, [administrator]},

%%
%% Additional network and protocol related configuration
%% =====================================================
%%

%% Set the default AMQP heartbeat delay (in seconds).
%%
%% {heartbeat, 600},

%% Set the max permissible size of an AMQP frame (in bytes).
%%
%% {frame_max, 131072},

%% Customising Socket Options.
%%
%% See (http://www.erlang.org/doc/man/inet.html#setopts-2) for
%% further documentation.
%%
%% {tcp_listen_options, [binary,
%%                       {packet,       raw},
%%                       {reuseaddr,    true},
%%                       {backlog,      128},
%%                       {nodelay,      true},
%%                       {exit_on_close, false}]},

%%
%% Resource Limits & Flow Control
%% ==============================
%%
%% See http://www.rabbitmq.com/memory.html for full details.

%% Memory-based Flow Control threshold.
%%
```

```
   %% {vm_memory_high_watermark, 0.4},

   %% Fraction of the high watermark limit at which queues start to
   %% page message out to disc in order to free up memory.
   %%
   %% {vm_memory_high_watermark_paging_ratio, 0.5},

   %% Set disk free limit (in bytes). Once free disk space reaches this
   %% lower bound, a disk alarm will be set - see the documentation
   %% listed above for more details.
   %%
   %% {disk_free_limit, 50000000},

   %% Alternatively, we can set a limit relative to total available RAM.
   %%
   %% {disk_free_limit, {mem_relative, 1.0}},

   %%
   %% Misc/Advanced Options
   %% ====================
   %%
   %% NB: Change these only if you understand what you are doing!
   %%

   %% To announce custom properties to clients on connection:
   %%
   %% {server_properties, []},

   %% How to respond to cluster partitions.
   %% See http://www.rabbitmq.com/partitions.html for further details.
   %%
   %% {cluster_partition_handling, ignore},

   %% Make clustering happen *automatically* at startup - only applied
   %% to nodes that have just been reset or started for the first time.
   %% See http://www.rabbitmq.com/clustering.html#auto-config for
   %% further details.
   %%
   %% {cluster_nodes, {['rabbit@my.host.com'], disc}},

   %% Set (internal) statistics collection granularity.
   %%
   %% {collect_statistics, none},

   %% Statistics collection interval (in milliseconds).
   %%
   %% {collect_statistics_interval, 5000},

   %% Explicitly enable/disable hipe compilation.
   %%
   %% {hipe_compile, true}

  ]},

 %% -----------------------------------------------------------------------
 %% Advanced Erlang Networking/Clustering Options.
 %%
 %% See http://www.rabbitmq.com/clustering.html for details
 %% -----------------------------------------------------------------------
 {kernel,
  [%% Provide an explicit port-range for inter-node communications.
   %% See http://www.rabbitmq.com/clustering.html#firewall for further details.

   %% Sets the minimum / maximum port numbers
   %%
   %% {inet_dist_listen_min, 10000},
   %% {inet_dist_listen_max, 10005},

   %% Sets the net_kernel tick time.
   %% Please see http://erlang.org/doc/man/kernel_app.html and
   %% http://www.rabbitmq.com/nettick.html for further details.
   %%
   %% {net_ticktime, 60}
```

```
 ]},

%% ----------------------------------------------------------------------
%% RabbitMQ Management Plugin
%%
%% See http://www.rabbitmq.com/management.html for details
%% ----------------------------------------------------------------------

{rabbitmq_management,
 [%% Pre-Load schema definitions from the following JSON file. See
  %% http://www.rabbitmq.com/management.html#load-definitions
  %%
  %% {load_definitions, "/path/to/schema.json"},

  %% Log all requests to the management HTTP API to a file.
  %%
  %% {http_log_dir, "/path/to/access.log"},

  %% Change the port on which the HTTP listener listens,
  %% specifying an interface for the web server to bind to.
  %% Also set the listener to use SSL and provide SSL options.
  %%
   {listener, [{port,     15672}
  %%             {ip,       "127.0.0.1"},
  %%             {ssl,      true},
  %%             {ssl_opts, [{cacertfile, "/path/to/cacert.pem"},
  %%                         {certfile,   "/path/to/cert.pem"},
  %%                         {keyfile,    "/path/to/key.pem"}]},
    ]}

  %% Configure how long aggregated data (such as message rates and queue
  %% lengths) is retained. Please read the plugin's documentation in
  %% https://www.rabbitmq.com/management.html#configuration for more
  %% details.
  %%
  %% {sample_retention_policies,
  %%  [{global,   [{60, 5}, {3600, 60}, {86400, 1200}]},
  %%   {basic,    [{60, 5}, {3600, 60}]},
  %%   {detailed, [{10, 5}]}]}
 ]},

{rabbitmq_management_agent,
 [%% Misc/Advanced Options
  %%
  %% NB: Change these only if you understand what you are doing!
  %%
  %% {force_fine_statistics, true}
 ]},

%% ----------------------------------------------------------------------
%% RabbitMQ Shovel Plugin
%%
%% See http://www.rabbitmq.com/shovel.html for details
%% ----------------------------------------------------------------------

{rabbitmq_shovel,
 [{shovels,
   [%% A named shovel worker.
    %% {my_first_shovel,
    %%  [

    %% List the source broker(s) from which to consume.
    %%
    %%   {sources,
    %%    [%% URI(s) and pre-declarations for all source broker(s).
    %%     {brokers, ["amqp://user:password@host.domain/my_vhost"]},
    %%     {declarations, []}
    %%    ]},

    %% List the destination broker(s) to publish to.
    %%   {destinations,
    %%    [%% A singular version of the 'brokers' element.
    %%     {broker, "amqp://"},
```

```
    %%     {declarations, []}
    %%    ]},

    %% Name of the queue to shovel messages from.
    %%
    %% {queue, <<"your-queue-name-goes-here">>},

    %% Optional prefetch count.
    %%
    %% {prefetch_count, 10},

    %% when to acknowledge messages:
    %% - no_ack: never (auto)
    %% - on_publish: after each message is republished
    %% - on_confirm: when the destination broker confirms receipt
    %%
    %% {ack_mode, on_confirm},

    %% Overwrite fields of the outbound basic.publish.
    %%
    %% {publish_fields, [{exchange,    <<"my_exchange">>},
    %%                   {routing_key, <<"from_shovel">>}]},

    %% Static list of basic.properties to set on re-publication.
    %%
    %% {publish_properties, [{delivery_mode, 2}]},

    %% The number of seconds to wait before attempting to
    %% reconnect in the event of a connection failure.
    %%
    {reconnect_delay, 2.5}

    %% ]} %% End of my_first_shovel
   ]}
  %% Rather than specifying some values per-shovel, you can specify
  %% them for all shovels here.
  %%
  %% {defaults, [{prefetch_count,     0},
  %%             {ack_mode,           on_confirm},
  %%             {publish_fields,     []},
  %%             {publish_properties, [{delivery_mode, 2}]},
  %%             {reconnect_delay,    2.5}]}
 ]},

%% -------------------------------------------------------------------
%% RabbitMQ Stomp Adapter
%%
%% See http://www.rabbitmq.com/stomp.html for details
%% -------------------------------------------------------------------

{rabbitmq_stomp,
 [%% Network Configuration - the format is generally the same as for the broker

  %% Listen only on localhost (ipv4 & ipv6) on a specific port.
  %% {tcp_listeners, [{"127.0.0.1", 61613},
  %%                  {"::1",       61613}]},

  %% Listen for SSL connections on a specific port.
  %% {ssl_listeners, [61614]},

  %% Additional SSL options

  %% Extract a name from the client's certificate when using SSL.
  %%
  %% {ssl_cert_login, true},

  %% Set a default user name and password. This is used as the default login
  %% whenever a CONNECT frame omits the login and passcode headers.
  %%
  %% Please note that setting this will allow clients to connect without
  %% authenticating!
  %%
  %% {default_user, [{login,    "guest"},
```

```
  %%              {passcode, "guest"}]},

  %% If a default user is configured, or you have configured use SSL client
  %% certificate based authentication, you can choose to allow clients to
  %% omit the CONNECT frame entirely. If set to true, the client is
  %% automatically connected as the default user or user supplied in the
  %% SSL certificate whenever the first frame sent on a session is not a
  %% CONNECT frame.
  %%
  %% {implicit_connect, true}
 ]},

%% -----------------------------------------------------------------------
%% RabbitMQ MQTT Adapter
%%
%% See http://hg.rabbitmq.com/rabbitmq-mqtt/file/stable/README.md for details
%% -----------------------------------------------------------------------

{rabbitmq_mqtt,
 [%% Set the default user name and password. Will be used as the default login
  %% if a connecting client provides no other login details.
  %%
  %% Please note that setting this will allow clients to connect without
  %% authenticating!
  %%
  %% {default_user, <<"guest">>},
  %% {default_pass, <<"guest">>},

  %% Enable anonymous access. If this is set to false, clients MUST provide
  %% login information in order to connect. See the default_user/default_pass
  %% configuration elements for managing logins without authentication.
  %%
  %% {allow_anonymous, true},

  %% If you have multiple chosts, specify the one to which the
  %% adapter connects.
  %%
  %% {vhost, <<"/">>},

  %% Specify the exchange to which messages from MQTT clients are published.
  %%
  %% {exchange, <<"amq.topic">>},

  %% Specify TTL (time to live) to control the lifetime of non-clean sessions.
  %%
  %% {subscription_ttl, 1800000},

  %% Set the prefetch count (governing the maximum number of unacknowledged
  %% messages that will be delivered).
  %%
  %% {prefetch, 10},

  %% TCP/SSL Configuration (as per the broker configuration).
  %%
  %% {tcp_listeners, [1883]},
  %% {ssl_listeners, []},

  %% TCP/Socket options (as per the broker configuration).
  %%
  %% {tcp_listen_options, [binary,
  %%                       {packet,    raw},
  %%                       {reuseaddr, true},
  %%                       {backlog,   128},
  %%                       {nodelay,   true}]}
 ]},

%% -----------------------------------------------------------------------
%% RabbitMQ AMQP 1.0 Support
%%
%% See http://hg.rabbitmq.com/rabbitmq-amqp1.0/file/default/README.md
%% for details
%% -----------------------------------------------------------------------
```

```
{rabbitmq_amqp1_0,
 [%% Connections that are not authenticated with SASL will connect as this
  %% account. See the README for more information.
  %%
  %% Please note that setting this will allow clients to connect without
  %% authenticating!
  %%
  %% {default_user, "guest"},

  %% Enable protocol strict mode. See the README for more information.
  %%
  %% {protocol_strict_mode, false}
 ]},

%% -----------------------------------------------------------------------
%% RabbitMQ LDAP Plugin
%%
%% See http://www.rabbitmq.com/ldap.html for details.
%%
%% -----------------------------------------------------------------------

{rabbitmq_auth_backend_ldap,
 [%%
  %% Connecting to the LDAP server(s)
  %% ================================
  %%

  %% Specify servers to bind to. You *must* set this in order for the plugin
  %% to work properly.
  %%
  %% {servers, ["your-server-name-goes-here"]},

  %% Connect to the LDAP server using SSL
  %%
  %% {use_ssl, false},

  %% Specify the LDAP port to connect to
  %%
  %% {port, 389},

  %% Enable logging of LDAP queries.
  %% One of
  %%   - false (no logging is performed)
  %%   - true (verbose logging of the logic used by the plugin)
  %%   - network (as true, but additionally logs LDAP network traffic)
  %%
  %% Defaults to false.
  %%
  %% {log, false},

  %%
  %% Authentication
  %% ==============
  %%

  %% Pattern to convert the username given through AMQP to a DN before
  %% binding
  %%
  %% {user_dn_pattern, "cn=${username},ou=People,dc=example,dc=com"},

  %% Alternatively, you can convert a username to a Distinguished
  %% Name via an LDAP lookup after binding. See the documentation for
  %% full details.

  %% When converting a username to a dn via a lookup, set these to
  %% the name of the attribute that represents the user name, and the
  %% base DN for the lookup query.
  %%
  %% {dn_lookup_attribute,   "userPrincipalName"},
  %% {dn_lookup_base,        "DC=gopivotal,DC=com"},

  %% Controls how to bind for authorisation queries and also to
  %% retrieve the details of users logging in without presenting a
```

	%% password (e.g., SASL EXTERNAL). %% One of %% - as_user (to bind as the authenticated user - requires a password) %% - anon (to bind anonymously) %% - {UserDN, Password} (to bind with a specified user name and password) %% %% Defaults to 'as_user'. %% %% {other_bind, as_user}, %% %% Authorisation %% ========== %% %% The LDAP plugin can perform a variety of queries against your %% LDAP server to determine questions of authorisation. See %% http://www.rabbitmq.com/ldap.html#authorisation for more %% information. %% Set the query to use when determining vhost access %% %% {vhost_access_query, {in_group, %% "ou=${vhost}-users,ou=vhosts,dc=example,dc=com"}}, %% Set the query to use when determining resource (e.g., queue) access %% %% {resource_access_query, {constant, true}}, %% Set queries to determine which tags a user has %% %% {tag_queries, []}]}].
Zawartość pliku konfiguracyjnego #2	/etc/rabbitmq/rabitmq-env.config ## # RabbitMQ environment variable names have the prefix RABBITMQ_.. # A typical variable called RABBITMQ_var_name is set as follows: # a shell environment variable called RABBITMQ_var_name is used. # if this is defined; # otherwise, a variable called var_name is used if this is set. # in the rabbitmq-env.conf file; # otherwise, a system-specified default value is used. ## ## # Default:. # the empty string - meaning bind to all network interfaces..<-> # Description:. # Change this if you only want to bind to one network interface. ## NODE_IP_ADDRESS="" ## # Default: # 5672 ## NODE_PORT=5672<> ## # Description: # The node name should be unique per erlang-node-and-machine combination. # To run multiple nodes, see the clustering guide. # Default: # Unix*: rabbit@$HOSTNAME # Windows: rabbit@%COMPUTERNAME% ## # NODENAME= ## # Description:

```
# The name of the installed service. This will appear in services.msc.
# Default:
# Windows Service: RabbitMQ
##############################################################################
# SERVICENAME

##############################################################################
# Description:
# Set this variable to new or reuse to redirect console output from the server
# to a file named %RABBITMQ_SERVICENAME%.debug in the default RABBITMQ_BASE.
# directory.
# If not set, console output from the server will be discarded (default).
# new A new file will be created each time the service starts.
# reuse The file will be overwritten each time the service starts.
# Default:
# Windows Service:
##############################################################################
# CONSOLE_LOG

##############################################################################
# Description:
# Parameters for the erl command used when invoking rabbitmqctl..
# This should be overridden for debugging purposes only.
# Default:
# None.>
##############################################################################
# CTL_ERL_ARGS

##############################################################################
# Desription:
# Standard parameters for the erl command used when invoking the RabbitMQ Server..
# This should be overridden for debugging purposes only.
# Default:
# Unix*:.
# "+K true +A30 +P 1048576 -kernel inet_default_connect_options [{nodelay,true}]"
# Windows: None
##############################################################################
# SERVER_ERL_ARGS

##############################################################################
# Description:
# Extra parameters for the erl command used when invoking the RabbitMQ Server..
# This will not override SERVER_ERL_ARGS..
# Default:
# None
##############################################################################
# SERVER_START_ARGS


##############################################################################
# File locations
##############################################################################
CONFIG_FILE=/etc/rabbitmq/rabbitmq

MNESIA_BASE=/opt/rabbitmq/mnesia
MNESIA_DIR=$MNESIA_BASE/$NODENAME

LOG_BASE=/opt/rabbitmq/log
LOGS=$LOG_BASE/$NODENAME.log
SASL_LOGS=$LOG_BASE/$NODENAME-sasl.log

PLUGINS_DIR=/usr/lib/rabbitmq/lib/rabbitmq_server-3.2.4/plugins
PLUGINS_EXPAND_DIR=$MNESIA_BASE/$NODENAME-plugins-expand
ENABLED_PLUGINS_FILE=/etc/rabbitmq/enabled_plugins

PID_FILE=$MNESIA_DIR/$NODENAME.pid
```

I. Nazwa serwera: szkzd1

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	szkzd1
IP / Domena	10.90.54.71
OS	CentOS 7
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	100 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Zarządzanie katalogiem

2. Konfiguracja .hosts

```
127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6

10.90.54.31 rt.epo.pl www.rt.epo.pl rt.epodreczniki.pl www.rt.epodrecnziki.pl epodreczniki.pcss.pl

10.90.54.200 epo.pl www.epo.pl epodreczniki.pl www.epodreczniki.pl
10.90.54.200 api.epo.pl www.api.epo.pl www.api.epodreczniki.pl
10.90.54.200 dev.epo.pl www.dev.epo.pl www.dev.epodreczniki.pl
10.90.54.200 static.epo.pl www.static.epo.pl static.epodreczniki.pl www.static.epodreczniki.pl
10.90.54.200 content.epo.pl www.content.epo.pl content.epodreczniki.pl www.content.epodreczniki.pl
10.90.54.200 preview.epo.pl www.preview.epo.pl preview.epodreczniki.pl www.preview.epodreczniki.pl
10.90.54.200 edit.epo.pl www.edit.epo.pl edit.epodreczniki.pl www.edit.epodreczniki.pl
10.90.54.200 repo.epo.pl www.repo.epo.pl repo.epodreczniki.pl www.repo.epodreczniki.pl
10.90.54.200 user.epo.pl www.user.epo.pl user.epodreczniki.pl www.user.epodreczniki.pl
10.90.54.200 search.epo.pl www.search.epo.pl www.search.epodreczniki.pl www.search.epodreczniki.pl

10.90.54.250 av.epo.pl www.av.epo.pl av.epodreczniki.pl www.av.epodreczniki.pl
```

3. Interfejsy sieciowe

eth0 10.90.54.71

4. Zainstalowane usługi

a. MAVEN

Parametr	Wartość
Lokalizacja instalacji	
Zawartość pliku konfiguracyjnego #1	/etc/maven/settings.xml <?xml version="1.0" encoding="UTF-8"?> <!-- Licensed to the Apache Software Foundation (ASF) under one or more contributor license agreements. See the NOTICE file distributed with this work for additional information regarding copyright ownership. The ASF licenses this file to you under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at http://www.apache.org/licenses/LICENSE-2.0

```
Unless required by applicable law or agreed to in writing,
software distributed under the License is distributed on an
"AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY
KIND, either express or implied.  See the License for the
specific language governing permissions and limitations
under the License.
-->

<!--
| This is the configuration file for Maven. It can be specified at two levels:
|
| 1. User Level. This settings.xml file provides configuration for a single user,.
|               and is normally provided in ${user.home}/.m2/settings.xml.
|
|               NOTE: This location can be overridden with the CLI option:
|
|               -s /path/to/user/settings.xml
|
| 2. Global Level. This settings.xml file provides configuration for all Maven
|               users on a machine (assuming they're all using the same Maven
|               installation). It's normally provided in.
|               ${maven.home}/conf/settings.xml.
|
|               NOTE: This location can be overridden with the CLI option:
|
|               -gs /path/to/global/settings.xml
|
| The sections in this sample file are intended to give you a running start at
| getting the most out of your Maven installation. Where appropriate, the default
| values (values used when the setting is not specified) are provided.
|
|-->
<settings xmlns="http://maven.apache.org/SETTINGS/1.0.0".
        xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance".
        xsi:schemaLocation="http://maven.apache.org/SETTINGS/1.0.0
http://maven.apache.org/xsd/settings-1.0.0.xsd">
  <!-- localRepository
   | The path to the local repository maven will use to store artifacts.
   |
   | Default: ~/.m2/repository
  <localRepository>/path/to/local/repo</localRepository>
  -->

  <!-- interactiveMode
   | This will determine whether maven prompts you when it needs input. If set to false,
   | maven will use a sensible default value, perhaps based on some other setting, for
   | the parameter in question.
   |
   | Default: true
  <interactiveMode>true</interactiveMode>
  -->

  <!-- offline
   | Determines whether maven should attempt to connect to the network when executing a build.
   | This will have an effect on artifact downloads, artifact deployment, and others.
   |
   | Default: false
  <offline>false</offline>
  -->

  <!-- pluginGroups
   | This is a list of additional group identifiers that will be searched when resolving plugins by their
prefix, i.e.
   | when invoking a command line like "mvn prefix:goal". Maven will automatically add the group
identifiers
   | "org.apache.maven.plugins" and "org.codehaus.mojo" if these are not already contained in the
list.
   |-->
  <pluginGroups>
    <!-- pluginGroup
     | Specifies a further group identifier to use for plugin lookup.
    <pluginGroup>com.your.plugins</pluginGroup>
    -->
```

```
  </pluginGroups>

  <!-- proxies
   | This is a list of proxies which can be used on this machine to connect to the network.
   | Unless otherwise specified (by system property or command-line switch), the first proxy
   | specification in this list marked as active will be used.
   |-->
  <proxies>
    <!-- proxy
     | Specification for one proxy, to be used in connecting to the network.
     |
    <proxy>
      <id>optional</id>
      <active>true</active>
      <protocol>http</protocol>
      <username>proxyuser</username>
      <password>proxypass</password>
      <host>proxy.host.net</host>
      <port>80</port>
      <nonProxyHosts>local.net|some.host.com</nonProxyHosts>
    </proxy>
    -->
  </proxies>

  <!-- servers
   | This is a list of authentication profiles, keyed by the server-id used within the system.
   | Authentication profiles can be used whenever maven must make a connection to a remote server.
   |-->
  <servers>
    <!-- server
     | Specifies the authentication information to use when connecting to a particular server, identified
by
     | a unique name within the system (referred to by the 'id' attribute below).
     |.
     | NOTE: You should either specify username/password OR privateKey/passphrase, since these
pairings are.
     |       used together.
     |
    <server>
      <id>deploymentRepo</id>
      <username>repouser</username>
      <password>repopwd</password>
    </server>
    -->
....
    <!-- Another sample, using keys to authenticate.
    <server>
      <id>siteServer</id>
      <privateKey>/path/to/private/key</privateKey>
      <passphrase>optional; leave empty if not used.</passphrase>
    </server>
    -->
  </servers>

  <!-- mirrors
   | This is a list of mirrors to be used in downloading artifacts from remote repositories.
   |.
   | It works like this: a POM may declare a repository to use in resolving certain artifacts.
   | However, this repository may have problems with heavy traffic at times, so people have mirrored
   | it to several places.
   |
   | That repository definition will have a unique id, so we can create a mirror reference for that
   | repository, to be used as an alternate download site. The mirror site will be the preferred.
   | server for that repository.
   |-->
  <mirrors>
    <!-- mirror
     | Specifies a repository mirror site to use instead of a given repository. The repository that
     | this mirror serves has an ID that matches the mirrorOf element of this mirror. IDs are used
     | for inheritance and direct lookup purposes, and must be unique across the set of mirrors.
     |
    <mirror>
      <id>mirrorId</id>
```

```
      <mirrorOf>repositoryId</mirrorOf>
      <name>Human Readable Name for this Mirror.</name>
      <url>http://my.repository.com/repo/path</url>
    </mirror>
    -->
  </mirrors>
..
  <!-- profiles
   | This is a list of profiles which can be activated in a variety of ways, and which can modify
   | the build process. Profiles provided in the settings.xml are intended to provide local machine-
   | specific paths and repository locations which allow the build to work in the local environment.
   |
   | For example, if you have an integration testing plugin - like cactus - that needs to know where
   | your Tomcat instance is installed, you can provide a variable here such that the variable is.
   | dereferenced during the build process to configure the cactus plugin.
   |
   | As noted above, profiles can be activated in a variety of ways. One way - the activeProfiles
   | section of this document (settings.xml) - will be discussed later. Another way essentially
   | relies on the detection of a system property, either matching a particular value for the property,
   | or merely testing its existence. Profiles can also be activated by JDK version prefix, where a.
   | value of '1.4' might activate a profile when the build is executed on a JDK version of '1.4.2_07'.
   | Finally, the list of active profiles can be specified directly from the command line.
   |
   | NOTE: For profiles defined in the settings.xml, you are restricted to specifying only artifact
   |       repositories, plugin repositories, and free-form properties to be used as configuration
   |       variables for plugins in the POM.
   |
   |-->
  <profiles>
    <!-- profile
     | Specifies a set of introductions to the build process, to be activated using one or more of the
     | mechanisms described above. For inheritance purposes, and to activate profiles via
<activatedProfiles/>
     | or the command line, profiles have to have an ID that is unique.
     |
     | An encouraged best practice for profile identification is to use a consistent naming convention
     | for profiles, such as 'env-dev', 'env-test', 'env-production', 'user-jdcasey', 'user-brett', etc.
     | This will make it more intuitive to understand what the set of introduced profiles is attempting
     | to accomplish, particularly when you only have a list of profile id's for debug.
     |
     | This profile example uses the JDK version to trigger activation, and provides a JDK-specific
repo.
    <profile>
      <id>jdk-1.4</id>

      <activation>
        <jdk>1.4</jdk>
      </activation>

      <repositories>
        <repository>
          <id>jdk14</id>
          <name>Repository for JDK 1.4 builds</name>
          <url>http://www.myhost.com/maven/jdk14</url>
          <layout>default</layout>
          <snapshotPolicy>always</snapshotPolicy>
        </repository>
      </repositories>
    </profile>
    -->

    <!--
     | Here is another profile, activated by the system property 'target-env' with a value of 'dev',
     | which provides a specific path to the Tomcat instance. To use this, your plugin configuration
     | might hypothetically look like:
     |
     | ...
     | <plugin>
     |   <groupId>org.myco.myplugins</groupId>
     |   <artifactId>myplugin</artifactId>
     |...
     |   <configuration>
     |     <tomcatLocation>${tomcatPath}</tomcatLocation>
```

		</configuration> | </plugin> | ...
NOTE: If you just wanted to inject this configuration whenever someone set 'target-env' to 	anything, you could just leave off the <value/> inside the activation-property. 	 <profile> <id>env-dev</id> <activation> <property> <name>target-env</name> <value>dev</value> </property> </activation> <properties> <tomcatPath>/path/to/tomcat/instance</tomcatPath> </properties> </profile> --> </profiles> <!-- activeProfiles | List of profiles that are active for all builds. | <activeProfiles> <activeProfile>alwaysActiveProfile</activeProfile> <activeProfile>anotherAlwaysActiveProfile</activeProfile> </activeProfiles> --> </settings>
Zawartość pliku konfiguracyjnego #2	<wklej treść pliku>

b. Ngnix

Parametr	Wartość
Lokalizacja instalacji	/etc/nginx

Zawartość pliku konfiguracyjnego #1	# For more information on configuration, see:
	# * Official English Documentation: http://nginx.org/en/docs/
	# * Official Russian Documentation: http://nginx.org/ru/docs/
	
	user nginx;
	worker_processes auto;
	error_log /var/log/nginx/error.log;
	pid /run/nginx.pid;
	
	# Load dynamic modules. See /usr/share/nginx/README.dynamic.
	include /usr/share/nginx/modules/*.conf;
	
	events {
	worker_connections 1024;
	}
	
	http {
	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;
	tcp_nodelay on;

	keepalive_timeout 65;
	types_hash_max_size 2048;
	
	include /etc/nginx/mime.types;
	default_type application/octet-stream;
	
	# Load modular configuration files from the /etc/nginx/conf.d directory.
	# See http://nginx.org/en/docs/ngx_core_module.html#include
	# for more information.
	include /etc/nginx/conf.d/*.conf;
	
	server {
	listen 80 default_server;
	listen [::]:80 default_server;
	server_name _;
	root /usr/share/nginx/html;
	
	# Load configuration files for the default server block.
	include /etc/nginx/default.d/*.conf;
	
	location / {
	}
	
	error_page 404 /404.html;
	location = /40x.html {
	}
	
	error_page 500 502 503 504 /50x.html;
	location = /50x.html {
	}
	}
	
	# Settings for a TLS enabled server.
	#
	# server {
	# listen 443 ssl http2 default_server;
	# listen [::]:443 ssl http2 default_server;
	# server_name _;
	# root /usr/share/nginx/html;
	#
	# ssl_certificate "/etc/pki/nginx/server.crt";
	# ssl_certificate_key "/etc/pki/nginx/private/server.key";
	# ssl_session_cache shared:SSL:1m;
	# ssl_session_timeout 10m;
	# ssl_ciphers HIGH:!aNULL:!MD5;
	# ssl_prefer_server_ciphers on;
	#
	# # Load configuration files for the default server block.
	# include /etc/nginx/default.d/*.conf;
	#
	# location / {
	# }
	#
	# error_page 404 /404.html;
	# location = /40x.html {
	# }
	#
	# error_page 500 502 503 504 /50x.html;
	# location = /50x.html {
	# }
	# }
	
	}

Zawartość pliku konfiguracyjnego #2	<wklej treść pliku>

c. JAVA

Parametr	Wartość
Lokalizacja instalacji	/etc/java/java.conf

Zawartość pliku konfiguracyjnego #1	# System-wide Java configuration file -*- sh -*-
	
	# Location of jar files on the system
	JAVA_LIBDIR=/usr/share/java
	
	# Location of arch-specific jar files on the system
	JNI_LIBDIR=/usr/lib/java
	
	# Root of all JVM installations
	JVM_ROOT=/usr/lib/jvm
	
	# You can define a system-wide JVM root here if you're not using the
	# default one.
	#
	# If you have a base JRE package installed
	# (e.g. java-1.6.0-openjdk):
	#JAVA_HOME=$JVM_ROOT/jre
	#
	# If you have a devel JDK package installed
	# (e.g. java-1.6.0-openjdk-devel):
	#JAVA_HOME=$JVM_ROOT/java
	
	# Options to pass to the java interpreter
	#JAVACMD_OPTS=
	
	# You can disable ABRT Java Connector by setting JAVA_ABRT to "off".
	# See: https://github.com/jfilak/abrt-java-connector/
	#JAVA_ABRT=off

I. Nazwa serwera: uwt1_elastic

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	uwt1_elastic
IP / Domena	10.90.54.111
OS	Ubuntu 16.04
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	200 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Wyszukiwanie treści

2. Konfiguracja .hosts

```
127.0.0.1 localhost

10.90.54.111    uwt1
10.90.54.112    uwt2

10.90.54.201    app1
10.90.54.202    app2
10.90.54.209    appcelery

10.90.54.211    storage1
10.90.54.212    storage2

10.90.54.221    db1

# internal proxy
10.90.54.200   epodreczniki.pl www.epodreczniki.pl api.epodreczniki.pl preview.epodreczniki.pl static.epodreczniki.pl
10.90.54.200   repo.epodreczniki.pl content.epodreczniki.pl user.epodreczniki.pl

# The following lines are desirable for IPv6 capable hosts
::1 ip6-localhost ip6-loopback
fe00::0 ip6-localnet
ff00::0 ip6-mcastprefix
ff02::1 ip6-allnodes
ff02::2 ip6-allrouters
ff02::3 ip6-allhosts
```

3. Interfejsy sieciowe

ens3 10.90.54.111

4. Zainstalowane usługi

a. Apache

Parametr	Wartość
Lokalizacja instalacji	/opt/elasitcsearch-1.7.3
Zawartość pliku konfiguracyjnego #1	/opt/elasitcsearch-1.7.3/config/elasticsearch.yml ##################### Elasticsearch Configuration Example ##################### # This file contains an overview of various configuration settings, # targeted at operations staff. Application developers should # consult the guide at <http://elasticsearch.org/guide>.

```
#
# The installation procedure is covered at
# <http://elasticsearch.org/guide/en/elasticsearch/reference/current/setup.html>.
#
# Elasticsearch comes with reasonable defaults for most settings,
# so you can try it out without bothering with configuration.
#
# Most of the time, these defaults are just fine for running a production
# cluster. If you're fine-tuning your cluster, or wondering about the
# effect of certain configuration option, please _do ask_ on the
# mailing list or IRC channel [http://elasticsearch.org/community].

# Any element in the configuration can be replaced with environment variables
# by placing them in ${...} notation. For example:
#
#node.rack: ${RACK_ENV_VAR}

# For information on supported formats and syntax for the config file, see
# <http://elasticsearch.org/guide/en/elasticsearch/reference/current/setup-configuration.html>


################################### Cluster ###################################

# Cluster name identifies your cluster for auto-discovery. If you're running
# multiple clusters on the same network, make sure you're using unique names.
#
#cluster.name: elasticsearch


#################################### Node #####################################

# Node names are generated dynamically on startup, so you're relieved
# from configuring them manually. You can tie this node to a specific name:
#
#node.name: "Franz Kafka"

# Every node can be configured to allow or deny being eligible as the master,
# and to allow or deny to store the data.
#
# Allow this node to be eligible as a master node (enabled by default):
#
#node.master: true
#
# Allow this node to store data (enabled by default):
#
#node.data: true

# You can exploit these settings to design advanced cluster topologies.
#
# 1. You want this node to never become a master node, only to hold data.
#    This will be the "workhorse" of your cluster.
#
#node.master: false
#node.data: true
#
# 2. You want this node to only serve as a master: to not store any data and
#    to have free resources. This will be the "coordinator" of your cluster.
#
#node.master: true
#node.data: false
#
# 3. You want this node to be neither master nor data node, but
#    to act as a "search load balancer" (fetching data from nodes,
#    aggregating results, etc.)
#
#node.master: false
#node.data: false

# Use the Cluster Health API [http://localhost:9200/_cluster/health], the
# Node Info API [http://localhost:9200/_nodes] or GUI tools
# such as <http://www.elasticsearch.org/overview/marvel/>,
# <http://github.com/karmi/elasticsearch-paramedic>,
# <http://github.com/lukas-vlcek/bigdesk> and
```

```
# <http://mobz.github.com/elasticsearch-head> to inspect the cluster state.

# A node can have generic attributes associated with it, which can later be used
# for customized shard allocation filtering, or allocation awareness. An attribute
# is a simple key value pair, similar to node.key: value, here is an example:
#
#node.rack: rack314

# By default, multiple nodes are allowed to start from the same installation location
# to disable it, set the following:
#node.max_local_storage_nodes: 1


################################### Index ###################################

# You can set a number of options (such as shard/replica options, mapping
# or analyzer definitions, translog settings, ...) for indices globally,
# in this file.
#
# Note, that it makes more sense to configure index settings specifically for
# a certain index, either when creating it or by using the index templates API.
#
# See <http://elasticsearch.org/guide/en/elasticsearch/reference/current/index-modules.html> and
# <http://elasticsearch.org/guide/en/elasticsearch/reference/current/indices-create-index.html>
# for more information.

# Set the number of shards (splits) of an index (5 by default):
#
#index.number_of_shards: 5

# Set the number of replicas (additional copies) of an index (1 by default):
#
#index.number_of_replicas: 1

# Note, that for development on a local machine, with small indices, it usually
# makes sense to "disable" the distributed features:
#
#index.number_of_shards: 1
#index.number_of_replicas: 0

# These settings directly affect the performance of index and search operations
# in your cluster. Assuming you have enough machines to hold shards and
# replicas, the rule of thumb is:
#
# 1. Having more *shards* enhances the _indexing_ performance and allows to
#    _distribute_ a big index across machines.
# 2. Having more *replicas* enhances the _search_ performance and improves the
#    cluster _availability_.
#
# The "number_of_shards" is a one-time setting for an index.
#
# The "number_of_replicas" can be increased or decreased anytime,
# by using the Index Update Settings API.
#
# Elasticsearch takes care about load balancing, relocating, gathering the
# results from nodes, etc. Experiment with different settings to fine-tune
# your setup.

# Use the Index Status API (<http://localhost:9200/A/_status>) to inspect
# the index status.


################################### Paths ###################################

# Path to directory containing configuration (this file and logging.yml):
#
#path.conf: /path/to/conf

# Path to directory where to store index data allocated for this node.
#
#path.data: /path/to/data
#
# Can optionally include more than one location, causing data to be striped across
```

```
# the locations (a la RAID 0) on a file level, favouring locations with most free
# space on creation. For example:
#
#path.data: /path/to/data1,/path/to/data2

# Path to temporary files:
#
#path.work: /path/to/work

# Path to log files:
#
#path.logs: /path/to/logs

# Path to where plugins are installed:
#
#path.plugins: /path/to/plugins


################################### Plugin ###################################

# If a plugin listed here is not installed for current node, the node will not start.
#
#plugin.mandatory: mapper-attachments,lang-groovy


################################### Memory ###################################

# Elasticsearch performs poorly when JVM starts swapping: you should ensure that
# it _never_ swaps.
#
#
# Set this property to true to lock the memory:
#
#bootstrap.mlockall: true

# Make sure that the ES_MIN_MEM and ES_MAX_MEM environment variables are set
# to the same value, and that the machine has enough memory to allocate
# for Elasticsearch, leaving enough memory for the operating system itself.
#
# You should also make sure that the Elasticsearch process is allowed to lock
# the memory, eg. by using `ulimit -l unlimited`.


############################## Network And HTTP ##############################

# Elasticsearch, by default, binds itself to the 0.0.0.0 address, and listens
# on port [9200-9300] for HTTP traffic and on port [9300-9400] for node-to-node
# communication. (the range means that if the port is busy, it will automatically
# try the next port).

# Set the bind address specifically (IPv4 or IPv6):
#
#network.bind_host: 192.168.0.1

# Set the address other nodes will use to communicate with this node. If not
# set, it is automatically derived. It must point to an actual IP address.
#
#network.publish_host: 192.168.0.1

# Set both 'bind_host' and 'publish_host':
#
#network.host: 192.168.0.1

# Set a custom port for the node to node communication (9300 by default):
#
#transport.tcp.port: 9300

# Enable compression for all communication between nodes (disabled by default):
#
#transport.tcp.compress: true

# Set a custom port to listen for HTTP traffic:
#
```

```
#http.port: 9200

# Set a custom allowed content length:
#
#http.max_content_length: 100mb

# Disable HTTP completely:
#
#http.enabled: false


################################### Gateway #####################################

# The gateway allows for persisting the cluster state between full cluster
# restarts. Every change to the state (such as adding an index) will be stored
# in the gateway, and when the cluster starts up for the first time,
# it will read its state from the gateway.

# There are several types of gateway implementations. For more information, see
# <http://elasticsearch.org/guide/en/elasticsearch/reference/current/modules-gateway.html>.

# The default gateway type is the "local" gateway (recommended):
#
#gateway.type: local

# Settings below control how and when to start the initial recovery process on
# a full cluster restart (to reuse as much local data as possible when using shared
# gateway).

# Allow recovery process after N nodes in a cluster are up:
#
#gateway.recover_after_nodes: 1

# Set the timeout to initiate the recovery process, once the N nodes
# from previous setting are up (accepts time value):
#
#gateway.recover_after_time: 5m

# Set how many nodes are expected in this cluster. Once these N nodes
# are up (and recover_after_nodes is met), begin recovery process immediately
# (without waiting for recover_after_time to expire):
#
#gateway.expected_nodes: 2


############################### Recovery Throttling #############################

# These settings allow to control the process of shards allocation between
# nodes during initial recovery, replica allocation, rebalancing,
# or when adding and removing nodes.

# Set the number of concurrent recoveries happening on a node:
#
# 1. During the initial recovery
#
#cluster.routing.allocation.node_initial_primaries_recoveries: 4
#
# 2. During adding/removing nodes, rebalancing, etc
#
#cluster.routing.allocation.node_concurrent_recoveries: 2

# Set to throttle throughput when recovering (eg. 100mb, by default 20mb):
#
#indices.recovery.max_bytes_per_sec: 20mb

# Set to limit the number of open concurrent streams when
# recovering a shard from a peer:
#
#indices.recovery.concurrent_streams: 5


################################### Discovery ###################################
```

```
# Discovery infrastructure ensures nodes can be found within a cluster
# and master node is elected. Multicast discovery is the default.

# Set to ensure a node sees N other master eligible nodes to be considered
# operational within the cluster. This should be set to a quorum/majority of.
# the master-eligible nodes in the cluster.
#
#discovery.zen.minimum_master_nodes: 1

# Set the time to wait for ping responses from other nodes when discovering.
# Set this option to a higher value on a slow or congested network
# to minimize discovery failures:
#
#discovery.zen.ping.timeout: 3s

# For more information, see
# <http://elasticsearch.org/guide/en/elasticsearch/reference/current/modules-discovery-zen.html>

# Unicast discovery allows to explicitly control which nodes will be used
# to discover the cluster. It can be used when multicast is not present,
# or to restrict the cluster communication-wise.
#
# 1. Disable multicast discovery (enabled by default):
#
#discovery.zen.ping.multicast.enabled: false
#
# 2. Configure an initial list of master nodes in the cluster
#    to perform discovery when new nodes (master or data) are started:
#
#discovery.zen.ping.unicast.hosts: ["host1", "host2:port"]

# EC2 discovery allows to use AWS EC2 API in order to perform discovery.
#
# You have to install the cloud-aws plugin for enabling the EC2 discovery.
#
# For more information, see
# <http://elasticsearch.org/guide/en/elasticsearch/reference/current/modules-discovery-ec2.html>
#
# See <http://elasticsearch.org/tutorials/elasticsearch-on-ec2/>
# for a step-by-step tutorial.

# GCE discovery allows to use Google Compute Engine API in order to perform discovery.
#
# You have to install the cloud-gce plugin for enabling the GCE discovery.
#
# For more information, see <https://github.com/elasticsearch/elasticsearch-cloud-gce>.

# Azure discovery allows to use Azure API in order to perform discovery.
#
# You have to install the cloud-azure plugin for enabling the Azure discovery.
#
# For more information, see <https://github.com/elasticsearch/elasticsearch-cloud-azure>.

################################## Slow Log ##################################

# Shard level query and fetch threshold logging.

#index.search.slowlog.threshold.query.warn: 10s
#index.search.slowlog.threshold.query.info: 5s
#index.search.slowlog.threshold.query.debug: 2s
#index.search.slowlog.threshold.query.trace: 500ms

#index.search.slowlog.threshold.fetch.warn: 1s
#index.search.slowlog.threshold.fetch.info: 800ms
#index.search.slowlog.threshold.fetch.debug: 500ms

#index.search.slowlog.threshold.fetch.warn: 1s
#index.search.slowlog.threshold.fetch.info: 800ms
#index.search.slowlog.threshold.fetch.debug: 500ms
#index.search.slowlog.threshold.fetch.trace: 200ms

#index.indexing.slowlog.threshold.index.warn: 10s
#index.indexing.slowlog.threshold.index.info: 5s
```

	#index.indexing.slowlog.threshold.index.debug: 2s #index.indexing.slowlog.threshold.index.trace: 500ms ################################## GC Logging ################################ #monitor.jvm.gc.young.warn: 1000ms #monitor.jvm.gc.young.info: 700ms #monitor.jvm.gc.young.debug: 400ms #monitor.jvm.gc.old.warn: 10s #monitor.jvm.gc.old.info: 5s #monitor.jvm.gc.old.debug: 2s ################################## Security ################################ # Uncomment if you want to enable JSONP as a valid return transport on the # http server. With this enabled, it may pose a security risk, so disabling # it unless you need it is recommended (it is disabled by default). # #http.jsonp.enable: true
Zawartość pliku konfiguracyjnego #2	/opt/elasticsearch-1.7.3/config/logging.yml # you can override this using by setting a system property, for example -Des.logger.level=DEBUG es.logger.level: INFO rootLogger: ${es.logger.level}, console, file logger: # log action execution errors for easier debugging action: DEBUG # reduce the logging for aws, too much is logged under the default INFO com.amazonaws: WARN org.apache.http: INFO # gateway #gateway: DEBUG #index.gateway: DEBUG # peer shard recovery #indices.recovery: DEBUG # discovery #discovery: TRACE index.search.slowlog: TRACE, index_search_slow_log_file index.indexing.slowlog: TRACE, index_indexing_slow_log_file additivity: index.search.slowlog: false index.indexing.slowlog: false appender: console: type: console layout: type: consolePattern conversionPattern: "[%d{ISO8601}][%-5p][%-25c] %m%n" file: type: dailyRollingFile file: ${path.logs}/${cluster.name}.log datePattern: "'.'yyyy-MM-dd" layout: type: pattern conversionPattern: "[%d{ISO8601}][%-5p][%-25c] %m%n" # Use the following log4j-extras RollingFileAppender to enable gzip compression of log files.. # For more information see https://logging.apache.org/log4j/extras/apidocs/org/apache/log4j/rolling/RollingFileAppender.html #file: #type: extrasRollingFile #file: ${path.logs}/${cluster.name}.log #rollingPolicy: timeBased #rollingPolicy.FileNamePattern: ${path.logs}/${cluster.name}.log.%d{yyyy-MM-dd}.gz #layout:

```
    #type: pattern
    #conversionPattern: "[%d{ISO8601}][%-5p][%-25c] %m%n"

index_search_slow_log_file:
  type: dailyRollingFile
  file: ${path.logs}/${cluster.name}_index_search_slowlog.log
  datePattern: "'.'yyyy-MM-dd"
  layout:
    type: pattern
    conversionPattern: "[%d{ISO8601}][%-5p][%-25c] %m%n"

index_indexing_slow_log_file:
  type: dailyRollingFile
  file: ${path.logs}/${cluster.name}_index_indexing_slowlog.log
  datePattern: "'.'yyyy-MM-dd"
  layout:
    type: pattern
    conversionPattern: "[%d{ISO8601}][%-5p][%-25c] %m%n"
```

I. Nazwa serwera: uwt2_solr

1. Parametry techniczne (metryczka)

Parametr	Wartość
Nazwa	uwt2_solr
IP / Domena	10.90.54.112
OS	Ubuntu 16.04
RAM:	32 GB
vCPU	10 vCPU
Przestrzeń	200 GB
Zewnętrzne IP:	NIE
Funkcja serwera:	Wyszukiwanie treści

2. Konfiguracja .hosts

```
127.0.0.1 localhost

10.90.54.111    uwt1
10.90.54.112    uwt2

10.90.54.200    proxy1

10.90.54.201    app1
10.90.54.202    app2
10.90.54.209    appcelery

10.90.54.211    storage1
10.90.54.212    storage2

10.90.54.221    db1

# internal proxy
10.90.54.200   epodreczniki.pl www.epodreczniki.pl api.epodreczniki.pl preview.epodreczniki.pl static.epodreczniki.pl
10.90.54.200   repo.epodreczniki.pl content.epodreczniki.pl user.epodreczniki.pl

# The following lines are desirable for IPv6 capable hosts
::1 ip6-localhost ip6-loopback
fe00::0 ip6-localnet
ff00::0 ip6-mcastprefix
ff02::1 ip6-allnodes
ff02::2 ip6-allrouters
ff02::3 ip6-allhosts
```

3. Interfejsy sieciowe

ens3 10.90.54.112

4. Zainstalowane usługi

a. Solr

Parametr	Wartość
Lokalizacja instalacji	

Zawartość pliku konfiguracyjnego #1	./
	# $HOME/.solr.in.sh
	# /usr/share/solr
	# /usr/local/share/solr
	# /opt/solr
	#
	# Another option is to specify the full path to the include file in the
	# environment. For example:
	#
	# $ SOLR_INCLUDE=/path/to/solr.in.sh solr start
	#
	# Note: This is particularly handy for running multiple instances on a
	# single installation, or for quick tests.
	#
	# Finally, developers and enthusiasts who frequently run from an SVN
	# checkout, and do not want to locally modify bin/solr.in.sh, can put
	# a customized include file at ~/.solr.in.sh.
	#
	# If you would rather configure startup entirely from the environment, you
	# can disable the include by exporting an empty SOLR_INCLUDE, or by
	# ensuring that no include files exist in the aforementioned search list.
	
	SOLR_SCRIPT="$0"
	verbose=false
	THIS_OS=`uname -s`
	
	if hash jar 2>/dev/null ; then # hash returns true if jar is on the path
	  UNPACK_WAR_CMD="$(command -v jar) xf"
	elif hash unzip 2>/dev/null ; then # hash returns true if unzip is on the path
	  UNPACK_WAR_CMD="$(command -v unzip) -q"
	else
	  echo -e "This script requires extracting a WAR file with either the jar or unzip
	utility, please install these utilities or contact your administrator for assistance."
	  exit 1
	fi
	
	stop_all=false

```
# for now, we don't support running this script from cygwin due to problems
# like not having lsof, ps auxww, curl, and awkward directory handling
if [ "${THIS_OS:0:6}" == "CYGWIN" ]; then
  echo -e "This script does not support cygwin due to severe limitations and lack
of adherence\nto BASH standards, such as lack of lsof, curl, and ps
options.\n\nPlease use the native solr.cmd script on Windows!"
  exit 1
fi

# Resolve symlinks to this script
while [ -h "$SOLR_SCRIPT" ] ; do
  ls=`ls -ld "$SOLR_SCRIPT"`
  # Drop everything prior to ->
  link=`expr "$ls" : '.*-> \(.*\)$'`
  if expr "$link" : '/.*' > /dev/null; then
    SOLR_SCRIPT="$link"
  else
    SOLR_SCRIPT=`dirname "$SOLR_SCRIPT"`/"$link"
  fi
done

SOLR_TIP=`dirname "$SOLR_SCRIPT"`/..
SOLR_TIP=`cd "$SOLR_TIP"; pwd`
DEFAULT_SERVER_DIR=$SOLR_TIP/example

# If an include wasn't specified in the environment, then search for one...
if [ "x$SOLR_INCLUDE" == "x" ]; then
    # Locations (in order) to use when searching for an include file.
    for include in "`dirname "$0"`/solr.in.sh" \
                   "$HOME/.solr.in.sh" \
                   /usr/share/solr/solr.in.sh \
                   /usr/local/share/solr/solr.in.sh \
                   /opt/solr/solr.in.sh; do
        if [ -r "$include" ]; then
            . "$include"
            break
        fi
    done
elif [ -r "$SOLR_INCLUDE" ]; then
    . "$SOLR_INCLUDE"
fi

if [ "$SOLR_JAVA_HOME" != "" ]; then
  JAVA=$SOLR_JAVA_HOME/bin/java
elif [ -n "$JAVA_HOME" ]; then
  for java in "$JAVA_HOME"/bin/amd64/java "$JAVA_HOME"/bin/java; do
    if [ -x "$java" ]; then
```

```
      JAVA="$java"
      break
    fi
  done
else
  JAVA=java
fi

# test that Java exists and is executable on this server
$JAVA -version >/dev/null 2>&1 || { echo >&2 "Java is required to run Solr!
Please install Java 7 or 8 before running this script."; exit 1; }

function print_usage() {
  CMD="$1"
  ERROR_MSG="$2"

  if [ "$ERROR_MSG" != "" ]; then
    echo -e "\nERROR: $ERROR_MSG\n"
  fi

  if [ "$CMD" == "" ]; then
    echo ""
    echo "Usage: solr COMMAND OPTIONS"
    echo "       where COMMAND is one of: start, stop, restart, healthcheck"
    echo ""
    echo "  Standalone server example (start Solr running in the background on port
8984):"
    echo ""
    echo "    ./solr start -p 8984"
    echo ""
    echo "  SolrCloud example (start Solr running in SolrCloud mode using
localhost:2181 to connect to ZooKeeper, with 1g max heap size and remote Java
debug options enabled):"
    echo ""
    echo "    ./solr start -c -m 1g -z localhost:2181 -a \"-Xdebug -
Xrunjdwp:transport=dt_socket,server=y,suspend=n,address=1044\""
    echo ""
    echo "Pass -help after any COMMAND to see command-specific usage
information,"
    echo "  such as:    ./solr start -help or ./solr stop -help"
    echo ""
  elif [[ "$CMD" == "start" || "$CMD" == "restart" ]]; then
    echo ""
    echo "Usage: solr $CMD [-f] [-c] [-h hostname] [-p port] [-d directory] [-z
zkHost] [-m memory] [-e example] [-s solr.solr.home] [-a \"additional-options\"]
[-V]"
    echo ""
    echo "  -f           Start Solr in foreground; default starts Solr in the background"
```

```
    echo "                and sends stdout / stderr to solr-PORT-console.log"
    echo ""
    echo "  -c or -cloud  Start Solr in SolrCloud mode; if -z not supplied, an
embedded ZooKeeper"
    echo "                instance is started on Solr port+1000, such as 9983 if Solr is
bound to 8983"
    echo ""
    echo "  -h <host>     Specify the hostname for this Solr instance"
    echo ""
    echo "  -p <port>     Specify the port to start the Solr HTTP listener on; default
is 8983"
    echo "                The specified port (SOLR_PORT) will also be used to
determine the stop port"
    echo "                STOP_PORT=(\$SOLR_PORT-1000) and JMX RMI listen
port RMI_PORT=(1\$SOLR_PORT). "
    echo "                For instance, if you set -p 8985, then the STOP_PORT=7985
and RMI_PORT=18985"
    echo ""
    echo "  -d <dir>      Specify the Solr server directory; defaults to server"
    echo ""
    echo "  -z <zkHost>   ZooKeeper connection string; only used when running in
SolrCloud mode using -c"
    echo "                To launch an embedded ZooKeeper instance, don't pass this
parameter."
    echo ""
    echo "  -m <memory>   Sets the min (-Xms) and max (-Xmx) heap size for the
JVM, such as: -m 4g"
    echo "                results in: -Xms4g -Xmx4g; by default, this script sets the
heap size to 512m"
    echo ""
    echo "  -s <dir>      Sets the solr.solr.home system property; Solr will create
core directories under"
    echo "                this directory. This allows you to run multiple Solr instances
on the same host"
    echo "                while reusing the same server directory set using the -d
parameter. If set, the"
    echo "                specified directory should contain a solr.xml file. The default
value is example/solr."
    echo "                This parameter is ignored when running examples (-e), as the
solr.solr.home depends"
    echo "                on which example is run."
    echo ""
    echo "  -e <example>  Name of the example to run; available examples:"
    echo "      cloud:         SolrCloud example"
    echo "      default:       Solr default example"
    echo "      dih:           Data Import Handler"
    echo "      schemaless:    Schema-less example"
    echo "      multicore:     Multicore"
```

```
    echo ""
    echo "  -a           Additional parameters to pass to the JVM when starting Solr,
such as to setup"
    echo "               Java debug options. For example, to enable a Java debugger
to attach to the Solr JVM"
    echo "               you could pass: -a \"-
agentlib:jdwp=transport=dt_socket,server=y,suspend=n,address=18983\""
    echo "               In most cases, you should wrap the additional parameters in
double quotes."
    echo ""
    echo "  -noprompt     Don't prompt for input; accept all defaults when running
examples that accept user input"
    echo ""
    echo "  -V           Verbose messages from this script"
    echo ""
  elif [ "$CMD" == "stop" ]; then
    echo ""
    echo "Usage: solr stop [-k key] [-p port] [-V]"
    echo ""
    echo "  -k <key>      Stop key; default is solrrocks"
    echo ""
    echo "  -p <port>     Specify the port the Solr HTTP listener is bound to"
    echo ""
    echo "  -all          Find and stop all running Solr servers on this host"
    echo ""
    echo "  NOTE: To see if any Solr servers are running, do: solr -i"
    echo ""
  elif [ "$CMD" == "healthcheck" ]; then
    echo ""
    echo "Usage: solr healthcheck [-c collection] [-z zkHost]"
    echo ""
    echo "  -c <collection>  Collection to run healthcheck against."
    echo ""
    echo "  -z <zkHost>      ZooKeeper connection string; default is localhost:9983"
    echo ""
  fi
} # end print_usage

# used to show the script is still alive when waiting on work to complete
spinner()
{
  local pid=$1
  local delay=0.5
  local spinstr='|/-\'
  while [ "$(ps a | awk '{print $1}' | grep $pid)" ]; do
      local temp=${spinstr#?}
      printf " [%c]  " "$spinstr"
      local spinstr=$temp${spinstr%"$temp"}
```

```
    sleep $delay
    printf "\b\b\b\b\b\b"
  done
  printf "    \b\b\b\b"
}

# given a port, find the pid for a Solr process
function solr_pid_by_port() {
  THE_PORT="$1"
  if [ -e "$SOLR_TIP/bin/solr-$THE_PORT.pid" ]; then
    PID=`cat $SOLR_TIP/bin/solr-$THE_PORT.pid`
    CHECK_PID=`ps auxww | awk '{print $2}' | grep $PID | sort -r | tr -d ' '`
    if [ "$CHECK_PID" != "" ]; then
      local solrPID=$PID
    fi
  fi
  echo "$solrPID"
}

# extract the value of the -Djetty.port parameter from a running Solr process
function jetty_port() {
  SOLR_PID="$1"
  SOLR_PROC=`ps auxww | grep $SOLR_PID | grep start.jar | grep jetty.port`
  IFS=' ' read -a proc_args <<< "$SOLR_PROC"
  for arg in "${proc_args[@]}"
    do
      IFS='=' read -a pair <<< "$arg"
      if [ "${pair[0]}" == "-Djetty.port" ]; then
        local jetty_port="${pair[1]}"
        break
      fi
    done
  echo "$jetty_port"
} # end jetty_port func

# run a Solr command-line tool using the SolrCLI class;
# useful for doing cross-platform work from the command-line using Java
function run_tool() {

  # Extract the solr.war if it hasn't been done already (so we can access the
SolrCLI class)
  if [[ -e $DEFAULT_SERVER_DIR/webapps/solr.war && ! -d
"$DEFAULT_SERVER_DIR/solr-webapp/webapp" ]]; then
    (mkdir -p $DEFAULT_SERVER_DIR/solr-webapp/webapp && cd
$DEFAULT_SERVER_DIR/solr-webapp/webapp && $UNPACK_WAR_CMD
$DEFAULT_SERVER_DIR/webapps/solr.war)
  fi
```

```
  "$JAVA" -Dlog4j.configuration=file:$DEFAULT_SERVER_DIR/scripts/cloud-
scripts/log4j.properties \
   -classpath "$DEFAULT_SERVER_DIR/solr-webapp/webapp/WEB-
INF/lib/*:$DEFAULT_SERVER_DIR/lib/ext/*" \
   org.apache.solr.util.SolrCLI $*

} # end run_tool function

# get information about any Solr nodes running on this host
function get_info() {
 # first, see if Solr is running
 numSolrs=`find $SOLR_TIP/bin -name "solr-*.pid" -type f | wc -l | tr -d ' '`
 if [ "$numSolrs" != "0" ]; then
  echo -e "\nFound $numSolrs Solr nodes: "
  for PIDF in `find $SOLR_TIP/bin -name "solr-*.pid" -type f`
   do
    ID=`cat $PIDF`
    port=`jetty_port "$ID"`
    if [ "$port" != "" ]; then
     echo ""
     echo "Solr process $ID running on port $port"
     run_tool status -solr http://localhost:$port/solr
     echo ""
    fi
  done
 else
  # no pid files but check using ps just to be sure
  numSolrs=`ps auxww | grep java | grep start.jar | wc -l | sed -e 's/^[ \t]*//'`
  if [ "$numSolrs" != "0" ]; then
   echo -e "\nFound $numSolrs Solr nodes: "
   for ID in `ps auxww | grep java | grep start.jar | awk '{print $2}' | sort -r`
    do
     port=`jetty_port "$ID"`
     if [ "$port" != "" ]; then
      echo ""
      echo "Solr process $ID running on port $port"
      run_tool status -solr http://localhost:$port/solr
      echo ""
     fi
   done
  else
   echo -e "\nNo Solr nodes are running.\n"
  fi
 fi

} # end get_info

# tries to gracefully stop Solr using the Jetty
```

```
# stop command and if that fails, then uses kill -9
function stop_solr() {

  DIR="$1"
  SOLR_PORT="$2"
  STOP_PORT=`expr $SOLR_PORT - 1000`
  STOP_KEY="$3"
  SOLR_PID="$4"

  if [ "$SOLR_PID" != "" ]; then
    echo -e "Sending stop command to Solr running on port $SOLR_PORT ...
waiting 5 seconds to allow Jetty process $SOLR_PID to stop gracefully."
    $JAVA -jar $DIR/start.jar STOP.PORT=$STOP_PORT
STOP.KEY=$STOP_KEY --stop || true
    (sleep 5) &
    spinner $!
    rm -f $SOLR_TIP/bin/solr-$SOLR_PORT.pid
  else
    echo -e "No Solr nodes found to stop."
    exit 0
  fi

  CHECK_PID=`ps auxww | awk '{print $2}' | grep $SOLR_PID | sort -r | tr -d ' '`
  if [ "$CHECK_PID" != "" ]; then
    echo -e "Solr process $SOLR_PID is still running; forcefully killing it now."
    kill -9 $SOLR_PID
    echo "Killed process $SOLR_PID"
    rm -f $SOLR_TIP/bin/solr-$SOLR_PORT.pid
    sleep 1
  fi

  CHECK_PID=`ps auxww | awk '{print $2}' | grep $SOLR_PID | sort -r | tr -d ' '`
  if [ "$CHECK_PID" != "" ]; then
    echo "ERROR: Failed to kill previous Solr Java process $SOLR_PID ... script
fails."
    exit 1
  fi
} # end stop_solr

if [ $# -eq 1 ]; then
  case $1 in
    -help|-usage)
        print_usage ""
        exit
    ;;
    -info|-i)
        get_info
        exit
```

```
    ;;
  esac
fi

if [ $# -gt 0 ]; then
  # if first arg starts with a dash (and it's not -help or -info),
  # then assume they are starting Solr, such as: solr -f
  if [[ $1 == -* ]]; then
    SCRIPT_CMD="start"
  else
    SCRIPT_CMD=$1
    shift
  fi
else
  # no args - just show usage and exit
  print_usage ""
  exit
fi

# run a healthcheck and exit if requested
if [ "$SCRIPT_CMD" == "healthcheck" ]; then

  if [ $# -gt 0 ]; then
    while true; do
      case $1 in
        -c|-collection)
          if [[ "$2" == "" || "${2:0:1}" == "-" ]]; then
            print_usage "$SCRIPT_CMD" "Expected collection name but found $2
instead!"
            exit 1
          fi
          HEALTHCHECK_COLLECTION=$2
          shift 2
        ;;
        -z|-zkhost)
          if [[ "$2" == "" || "${2:0:1}" == "-" ]]; then
            print_usage "$SCRIPT_CMD" "Expected a ZooKeeper connection
string but found $2 instead!"
            exit 1
          fi
          ZK_HOST="$2"
          shift 2
        ;;
        -help|-usage)
          print_usage "$SCRIPT_CMD"
          exit 0
        ;;
        --)
```

```
            shift
            break
        ;;
        *)
            if [ "$1" != "" ]; then
              print_usage "$SCRIPT_CMD" "Unrecognized or misplaced argument:
$1!"
              exit 1
            else
              break # out-of-args, stop looping
            fi
        ;;
    esac
  done
 fi

 if [ "$ZK_HOST" == "" ]; then
   ZK_HOST=localhost:9983
 fi

 if [ "$HEALTHCHECK_COLLECTION" == "" ]; then
   echo "collection parameter is required!"
   print_usage "healthcheck"
   exit 1
 fi

 run_tool healthcheck -zkHost $ZK_HOST -collection
$HEALTHCHECK_COLLECTION

 exit $?
fi

# verify the command given is supported
if [ "$SCRIPT_CMD" != "stop" ] && [ "$SCRIPT_CMD" != "start" ] && [
"$SCRIPT_CMD" != "restart" ]; then
 print_usage "" "$SCRIPT_CMD not supported!"
 exit 1
fi

# Run in foreground (default is to run in the background)
FG="false"
noprompt=false

if [ $# -gt 0 ]; then
 while true; do
   case $1 in
     -c|-cloud)
       SOLR_MODE="solrcloud"
```

```
        shift
      ;;
      -d|-dir)
        if [[ "$2" == "" || "${2:0:1}" == "-" ]]; then
          print_usage "$SCRIPT_CMD" "Expected directory but found $2
instead!"
          exit 1
        fi

        if [[ "$2" == "." || "$2" == "./" || "$2" == ".." || "$2" == "../" ]]; then
          SOLR_SERVER_DIR=`pwd`/$2
        else
          # see if the arg value is relative to the tip vs full path
          if [[ $2 != /* ]] && [[ -d "$SOLR_TIP/$2" ]]; then
            SOLR_SERVER_DIR="$SOLR_TIP/$2"
          else
            SOLR_SERVER_DIR="$2"
          fi
        fi
        # resolve it to an absolute path
        SOLR_SERVER_DIR=`cd "$SOLR_SERVER_DIR"; pwd`
        shift 2
      ;;
      -s|-solr.home)
        if [[ "$2" == "" || "${2:0:1}" == "-" ]]; then
          print_usage "$SCRIPT_CMD" "Expected directory but found $2
instead!"
          exit 1
        fi

        SOLR_HOME="$2"
        shift 2
      ;;
      -e|-example)
        if [[ "$2" == "" || "${2:0:1}" == "-" ]]; then
          print_usage "$SCRIPT_CMD" "Expected example name but found $2
instead!"
          exit 1
        fi
        EXAMPLE="$2"
        shift 2
      ;;
      -f|-foreground)
        FG="true"
        shift
      ;;
      -h|-host)
        if [[ "$2" == "" || "${2:0:1}" == "-" ]]; then
```

```
            print_usage "$SCRIPT_CMD" "Expected hostname but found $2
instead!"
            exit 1
          fi
          SOLR_HOST="$2"
          shift 2
        ;;
        -m|-memory)
          if [[ "$2" == "" || "${2:0:1}" == "-" ]]; then
            print_usage "$SCRIPT_CMD" "Expected memory setting but found $2
instead!"
            exit 1
          fi
          SOLR_HEAP="$2"
          shift 2
        ;;
        -p|-port)
          if [[ "$2" == "" || "${2:0:1}" == "-" ]]; then
            print_usage "$SCRIPT_CMD" "Expected port number but found $2
instead!"
            exit 1
          fi
          SOLR_PORT="$2"
          shift 2
        ;;
        -z|-zkhost)
          if [[ "$2" == "" || "${2:0:1}" == "-" ]]; then
            print_usage "$SCRIPT_CMD" "Expected ZooKeeper connection string
but found $2 instead!"
            exit 1
          fi

          ZK_HOST="$2"
          shift 2
        ;;
        -a|-addlopts)
          ADDITIONAL_CMD_OPTS="$2"
          shift 2
        ;;
        -k|-key)
          STOP_KEY="$2"
          shift 2
        ;;
        -help|-usage)
          print_usage "$SCRIPT_CMD"
          exit 0
        ;;
        -noprompt)
```

```
        noprompt=true
        shift
      ;;
      -V|-verbose)
        verbose=true
        shift
      ;;
      -all)
        stop_all=true
        shift
      ;;
      --)
        shift
        break
      ;;
      *)
        if [ "$1" != "" ]; then
          print_usage "$SCRIPT_CMD" "Error parsing argument $1!"
          exit 1
        else
          break # out-of-args, stop looping
        fi
      ;;
    esac
  done
fi

if $verbose ; then
  echo "Using Solr root directory: $SOLR_TIP"
  echo "Using Java: $JAVA"
  $JAVA -version
fi

if [ "$SOLR_HOST" != "" ]; then
  SOLR_HOST_ARG="-Dhost=$SOLR_HOST"
else
  SOLR_HOST_ARG=""
fi

if [ "$SOLR_SERVER_DIR" == "" ]; then
  SOLR_SERVER_DIR=$DEFAULT_SERVER_DIR
fi

if [ ! -e "$SOLR_SERVER_DIR" ]; then
  echo -e "\nSolr server directory $SOLR_SERVER_DIR not found!\n"
  exit 1
fi
```

```
CLOUD_NUM_NODES=2
declare -a CLOUD_PORTS=('8983' '7574' '8984' '7575');

# select solr.solr.home based on the desired example
if [ "$EXAMPLE" != "" ]; then
   case $EXAMPLE in
      cloud)
        #
        # Engage in an interactive session with user to setup the SolrCloud
example
        #
        echo -e "\nWelcome to the SolrCloud example!\n\n"
        if $noprompt ; then
          CLOUD_NUM_NODES=2
          echo -e "Starting up $CLOUD_NUM_NODES Solr nodes for your
example SolrCloud cluster."
        else
          echo -e "This interactive session will help you launch a SolrCloud cluster
on your local workstation.\n"
          read -e -p "To begin, how many Solr nodes would you like to run in your
local cluster? (specify 1-4 nodes) [2] " USER_INPUT
          while true
          do
            CLOUD_NUM_NODES=`echo $USER_INPUT | tr -d ' '`
            if [ "$CLOUD_NUM_NODES" == "" ]; then
              CLOUD_NUM_NODES=2
            fi
            if [[ $CLOUD_NUM_NODES > 4 || $CLOUD_NUM_NODES < 1 ]];
then
              read -e -p "Please provide a node count between 1 and 4 [2] "
USER_INPUT
            else
              break;
            fi
          done

          echo -e "Ok, let's start up $CLOUD_NUM_NODES Solr nodes for your
example SolrCloud cluster.\n"
          for (( s=0; s<$CLOUD_NUM_NODES; s++ ))
          do
            read -e -p "Please enter the port for node$[$s+1]
[${CLOUD_PORTS[$s]}] " USER_INPUT
            while true
            do
              # trim whitespace out of the user input
              CLOUD_PORT=`echo $USER_INPUT | tr -d ' '`

              # handle the default selection or empty input
```

```
            if [ "$CLOUD_PORT" == "" ]; then
              CLOUD_PORT=${CLOUD_PORTS[$s]}
            fi

            # check to see if something is already bound to that port
            if hash lsof 2>/dev/null ; then  # hash returns true if lsof is on the path
              PORT_IN_USE=`lsof -Pni:$CLOUD_PORT`
              if [ "$PORT_IN_USE" != "" ]; then
                read -e -p "Oops! Looks like port $CLOUD_PORT is already being
used by another process. Please choose a different port. " USER_INPUT
              else
                CLOUD_PORTS[$s]=$CLOUD_PORT
                echo $CLOUD_PORT
                break;
              fi
            else
              CLOUD_PORTS[$s]=$CLOUD_PORT
              echo $CLOUD_PORT
              break;
            fi
          done
        done
      fi

      for (( s=0; s<$CLOUD_NUM_NODES; s++ ))
      do
        ndx=$[$s+1]
        if [ ! -d "$SOLR_TIP/node$ndx" ]; then
          echo "Cloning $DEFAULT_SERVER_DIR into
$SOLR_TIP/node$ndx"
          cp -r $DEFAULT_SERVER_DIR $SOLR_TIP/node$ndx
          rm -rf $SOLR_TIP/node$ndx/solr/zoo_data
        fi
      done
      SOLR_MODE="solrcloud"
      SOLR_SERVER_DIR="$SOLR_TIP/node1"
      SOLR_HOME="$SOLR_SERVER_DIR/solr"
      SOLR_PORT=${CLOUD_PORTS[0]}
      shift
    ;;
    default)
      SOLR_HOME="$SOLR_SERVER_DIR/solr"
      shift
    ;;
    dih)
      SOLR_HOME="$SOLR_SERVER_DIR/example-DIH/solr"
      shift
    ;;
```

```
      schemaless)
        SOLR_HOME="$SOLR_SERVER_DIR/example-schemaless/solr"
        shift
      ;;
      multicore)
        SOLR_HOME="$SOLR_SERVER_DIR/multicore"
        shift
      ;;
      *)
        print_usage "start" "Unsupported example $EXAMPLE! Please choose
one of: cloud, dih, schemaless, multicore, or default"
        exit 1
      ;;
  esac
fi

if [[ "$FG" == "true" && "$EXAMPLE" != "" ]]; then
  FG="false"
  echo -e "\nWARNING: Foreground mode (-f) not supported when running
examples.\n"
fi

if [ "$STOP_KEY" == "" ]; then
  STOP_KEY="solrrocks"
fi

# stop all if no port specified
if [[ "$SCRIPT_CMD" == "stop" && "$SOLR_PORT" == "" ]]; then
  if $stop_all; then
    none_stopped=true
    for PIDF in `find $SOLR_TIP/bin -name "solr-*.pid" -type f`
      do
        NEXT_PID=`cat $PIDF`
        port=`jetty_port "$NEXT_PID"`
        if [ "$port" != "" ]; then
          stop_solr "$SOLR_SERVER_DIR" "$port" "$STOP_KEY" "$NEXT_PID"
          none_stopped=false
        fi
        rm -f $PIDF
    done
    if $none_stopped; then
      echo -e "\nNo Solr nodes found to stop.\n"
    fi
  else
    echo -e "\nERROR: Must either specify a port using -p or -all to stop all Solr
nodes on this host.\n"
    exit 1
  fi
```

```
  exit
fi

if [ "$SOLR_PORT" == "" ]; then
  SOLR_PORT="8983"
fi

if [ "$STOP_PORT" == "" ]; then
  STOP_PORT=`expr $SOLR_PORT - 1000`
fi

if [[ "$SCRIPT_CMD" == "start" ]]; then
  # see if Solr is already running
  SOLR_PID=`solr_pid_by_port "$SOLR_PORT"`

  if [ "$SOLR_PID" == "" ]; then
    # not found using the pid file ... but use ps to ensure not found
    SOLR_PID=`ps auxww | grep start.jar | grep $SOLR_PORT | grep -v grep |
awk '{print $2}' | sort -r`
  fi

  if [ "$SOLR_PID" != "" ]; then
    echo -e "\nPort $SOLR_PORT is already being used by another process (pid:
$SOLR_PID)\n"
    exit 1
  fi
else
  # either stop or restart
  # see if Solr is already running
  SOLR_PID=`solr_pid_by_port "$SOLR_PORT"`
  if [ "$SOLR_PID" == "" ]; then
    # not found using the pid file ... but use ps to ensure not found
    SOLR_PID=`ps auxww | grep start.jar | grep $SOLR_PORT | grep -v grep |
awk '{print $2}' | sort -r`
  fi
  if [ "$SOLR_PID" != "" ]; then
    stop_solr "$SOLR_SERVER_DIR" "$SOLR_PORT" "$STOP_KEY"
"$SOLR_PID"
  else
    if [ "$SCRIPT_CMD" == "stop" ]; then
      echo -e "No process found for Solr node running on port $SOLR_PORT"
      exit 1
    fi
  fi
fi

# backup the log files
if [ -f $SOLR_SERVER_DIR/logs/solr.log ]; then
```

```
  if $verbose ; then
    echo "Backing up $SOLR_SERVER_DIR/logs/solr.log"
  fi
  mv $SOLR_SERVER_DIR/logs/solr.log
$SOLR_SERVER_DIR/logs/solr_log_`date +"%Y%m%d_%H%M"`
fi

if [ -f $SOLR_SERVER_DIR/logs/solr_gc.log ]; then
  if $verbose ; then
    echo "Backing up $SOLR_SERVER_DIR/logs/solr_gc.log"
  fi
  mv $SOLR_SERVER_DIR/logs/solr_gc.log
$SOLR_SERVER_DIR/logs/solr_gc_log_`date +"%Y%m%d_%H%M"`
fi

if [ "$SCRIPT_CMD" == "stop" ]; then
  # already stopped, script is done.
  exit 0
fi

# if we get here, then we're starting a new node up ...
if [ "$SOLR_HOME" == "" ]; then
  SOLR_HOME="$SOLR_SERVER_DIR/solr"
else
  if [[ $SOLR_HOME != /* ]] && [[ -d
"$SOLR_SERVER_DIR/$SOLR_HOME" ]]; then
    SOLR_HOME="$SOLR_SERVER_DIR/$SOLR_HOME"
  fi
fi

if [ ! -e "$SOLR_HOME" ]; then
  echo -e "\nSolr home directory $SOLR_HOME not found!\n"
  exit 1
fi
if [ ! -e "$SOLR_HOME/solr.xml" ]; then
  echo -e "\nSolr home directory $SOLR_HOME must contain a solr.xml file!\n"
  exit 1
fi

# if verbose gc logging enabled, setup the location of the log file
if [ "$GC_LOG_OPTS" != "" ]; then
  GC_LOG_OPTS="$GC_LOG_OPTS -
Xloggc:$SOLR_SERVER_DIR/logs/solr_gc.log"
fi

if [ "$SOLR_MODE" == "solrcloud" ]; then
  if [ "$ZK_CLIENT_TIMEOUT" == "" ]; then
    ZK_CLIENT_TIMEOUT="15000"
```

```
  fi

  CLOUD_MODE_OPTS="-DzkClientTimeout=$ZK_CLIENT_TIMEOUT"

  if [ "$ZK_HOST" != "" ]; then
    CLOUD_MODE_OPTS="$CLOUD_MODE_OPTS -DzkHost=$ZK_HOST"
  else
    if $verbose ; then
      echo "Configuring SolrCloud to launch an embedded ZooKeeper using -
DzkRun"
    fi

    CLOUD_MODE_OPTS="$CLOUD_MODE_OPTS -DzkRun"
  fi

  # and if collection1 needs to be bootstrapped
  if [ -e "$SOLR_HOME/collection1/core.properties" ]; then
    CLOUD_MODE_OPTS="$CLOUD_MODE_OPTS -
Dbootstrap_confdir=./solr/collection1/conf -Dcollection.configName=myconf -
DnumShards=1"
  fi

fi

# These are useful for attaching remove profilers like VisualVM/JConsole
if [ "$ENABLE_REMOTE_JMX_OPTS" == "true" ]; then

  if [ "$RMI_PORT" == "" ]; then
    RMI_PORT=1$SOLR_PORT
  fi

  REMOTE_JMX_OPTS="-Dcom.sun.management.jmxremote \
-Dcom.sun.management.jmxremote.local.only=false \
-Dcom.sun.management.jmxremote.ssl=false \
-Dcom.sun.management.jmxremote.authenticate=false \
-Dcom.sun.management.jmxremote.port=$RMI_PORT \
-Dcom.sun.management.jmxremote.rmi.port=$RMI_PORT"

  # if the host is set, then set that as the rmi server hostname
  if [ "$SOLR_HOST" != "" ]; then
    REMOTE_JMX_OPTS="$REMOTE_JMX_OPTS -
Djava.rmi.server.hostname=$SOLR_HOST"
  fi
else
  REMOTE_JMX_OPTS=""
fi

if [ "$SOLR_HEAP" != "" ]; then
```

```
  SOLR_JAVA_MEM="-Xms$SOLR_HEAP -Xmx$SOLR_HEAP"
fi

if [ "$SOLR_JAVA_MEM" == "" ]; then
  SOLR_JAVA_MEM="-Xms512m -Xmx512m"
fi

if [ "$SOLR_TIMEZONE" == "" ]; then
  SOLR_TIMEZONE="UTC"
fi

# Launches Solr in foreground/background depending on parameters
function launch_solr() {

  run_in_foreground="$1"
  stop_port="$STOP_PORT"

  SOLR_ADDL_ARGS="$2"

  # deal with Java version specific GC and other flags
  JAVA_VERSION=`echo "$($JAVA -version 2>&1)" | grep "java version" | awk
'{ print substr($3, 2, length($3)-2); }'`
  if [ "${JAVA_VERSION:0:3}" == "1.7" ]; then
    # Specific Java version hacking
    GC_TUNE="$GC_TUNE -XX:CMSFullGCsBeforeCompaction=1 -
XX:CMSTriggerPermRatio=80"
    JAVA_MINOR_VERSION=${JAVA_VERSION:(-2)}
    if [[ $JAVA_MINOR_VERSION -ge 40 && $JAVA_MINOR_VERSION -le
51 ]]; then
      GC_TUNE="$GC_TUNE -XX:-UseSuperWord"
      echo -e "\nWARNING: Java version $JAVA_VERSION has known bugs
with Lucene and requires the -XX:-UseSuperWord flag. Please consider
upgrading your JVM.\n"
    fi
  fi

  if $verbose ; then
    echo -e "\nStarting Solr using the following settings:"
    echo -e "    JAVA            = $JAVA"
    echo -e "    SOLR_SERVER_DIR = $SOLR_SERVER_DIR"
    echo -e "    SOLR_HOME       = $SOLR_HOME"
    echo -e "    SOLR_HOST       = $SOLR_HOST"
    echo -e "    SOLR_PORT       = $SOLR_PORT"
    echo -e "    STOP_PORT       = $STOP_PORT"
    echo -e "    SOLR_JAVA_MEM   = $SOLR_JAVA_MEM"
    echo -e "    GC_TUNE         = $GC_TUNE"
    echo -e "    GC_LOG_OPTS     = $GC_LOG_OPTS"
    echo -e "    SOLR_TIMEZONE   = $SOLR_TIMEZONE"
```

```
  if [ "$SOLR_MODE" == "solrcloud" ]; then
    echo -e "    CLOUD_MODE_OPTS = $CLOUD_MODE_OPTS"
  fi

  if [ "$SOLR_ADDL_ARGS" != "" ]; then
    echo -e "    SOLR_ADDL_ARGS   = $SOLR_ADDL_ARGS"
  fi

  if [ "$ENABLE_REMOTE_JMX_OPTS" == "true" ]; then
    echo -e "    RMI_PORT        = $RMI_PORT"
    echo -e "    REMOTE_JMX_OPTS = $REMOTE_JMX_OPTS"
  fi
  echo -e "\n"
 fi

 # need to launch solr from the server dir
 cd $SOLR_SERVER_DIR

 if [ ! -e "$SOLR_SERVER_DIR/start.jar" ]; then
   echo -e "\nERROR: start.jar file not found in $SOLR_SERVER_DIR!\nPlease
check your -d parameter to set the correct Solr server directory.\n"
   exit 1
 fi

 SOLR_START_OPTS="-server -Xss256k $SOLR_JAVA_MEM $GC_TUNE
$GC_LOG_OPTS $REMOTE_JMX_OPTS \
 $CLOUD_MODE_OPTS \
-DSTOP.PORT=$stop_port -DSTOP.KEY=$STOP_KEY \
$SOLR_HOST_ARG -Djetty.port=$SOLR_PORT \
-Dsolr.solr.home=$SOLR_HOME \
-Dsolr.install.dir=$SOLR_TIP \
-Duser.timezone=$SOLR_TIMEZONE \
-Djava.net.preferIPv4Stack=true"

 if [ "$SOLR_MODE" == "solrcloud" ]; then
   IN_CLOUD_MODE=" in SolrCloud mode"
 fi

 mkdir -p $SOLR_SERVER_DIR/logs

 if [ "$run_in_foreground" == "true" ]; then
   echo -e "\nStarting Solr$IN_CLOUD_MODE on port $SOLR_PORT from
$SOLR_SERVER_DIR\n"
   $JAVA $SOLR_START_OPTS $SOLR_ADDL_ARGS -
XX:OnOutOfMemoryError="$SOLR_TIP/bin/oom_solr.sh $SOLR_PORT" -jar
start.jar
 else
```

```
    # run Solr in the background
    nohup $JAVA $SOLR_START_OPTS $SOLR_ADDL_ARGS -
XX:OnOutOfMemoryError="$SOLR_TIP/bin/oom_solr.sh $SOLR_PORT" -jar
start.jar 1>$SOLR_SERVER_DIR/logs/solr-$SOLR_PORT-console.log 2>&1 &
echo $! > $SOLR_TIP/bin/solr-$SOLR_PORT.pid

    # no lsof on cygwin though
    if hash lsof 2>/dev/null ; then  # hash returns true if lsof is on the path
      echo -n "Waiting to see Solr listening on port $SOLR_PORT"
      # Launch in a subshell to show the spinner
      (loops=0
      while true
      do
        running=`lsof -Pni:$SOLR_PORT`
        if [ "$running" == "" ]; then
          if [ $loops -lt 6 ]; then
            sleep 5
            loops=$[$loops+1]
          else
            echo -e "Still not seeing Solr listening on $SOLR_PORT after 30
seconds!"
            tail -30 $SOLR_SERVER_DIR/logs/solr.log
            exit;
          fi
        else
          SOLR_PID=`ps auxww | grep start.jar | grep $SOLR_PORT | grep -v grep |
awk '{print $2}' | sort -r`
          echo -e "\nStarted Solr server on port $SOLR_PORT (pid=$SOLR_PID).
Happy searching!\n"
          exit;
        fi
      done) &
      spinner $!
    else
      SOLR_PID=`ps auxww | grep start.jar | grep $SOLR_PORT | grep -v grep |
awk '{print $2}' | sort -r`
      echo -e "\nStarted Solr server on port $SOLR_PORT (pid=$SOLR_PID).
Happy searching!\n"
      exit;
    fi
  fi
}

if [ "$EXAMPLE" != "cloud" ]; then
  launch_solr "$FG" "$ADDITIONAL_CMD_OPTS"
else
  #
  # SolrCloud example is a bit involved so needs special handling here
```

```
  #
  SOLR_SERVER_DIR=$SOLR_TIP/node1
  SOLR_HOME=$SOLR_TIP/node1/solr
  SOLR_PORT=${CLOUD_PORTS[0]}

  if [ "$ZK_HOST" != "" ]; then
    DASHZ="-z $ZK_HOST"
  fi

  if [ "$SOLR_HEAP" != "" ]; then
    DASHM="-m $SOLR_HEAP"
  fi

  if [ "$ADDITIONAL_CMD_OPTS" != "" ]; then
    DASHA="-a $ADDITIONAL_CMD_OPTS"
  fi

  echo -e "\nStarting up SolrCloud node1 on port ${CLOUD_PORTS[0]} using
command:\n"
  echo -e "solr start -cloud -d node1 -p $SOLR_PORT $DASHZ $DASHM
$DASHA\n\n"

  # can't launch this node in the foreground else we can't run anymore commands
  launch_solr "false" "$ADDITIONAL_CMD_OPTS"

  sleep 5

  # if user did not define a specific -z parameter, assume embedded in first cloud
node we launched above
  zk_host=$ZK_HOST
  if [ "$zk_host" == "" ]; then
    zk_port=$[$SOLR_PORT+1000]
    zk_host=localhost:$zk_port
  fi

  for (( s=1; s<$CLOUD_NUM_NODES; s++ ))
  do
    ndx=$[$s+1]
    next_port=${CLOUD_PORTS[$s]}
    echo -e "\n\nStarting node$ndx on port $next_port using command:\n"
    echo -e "solr start -cloud -d node$ndx -p $next_port -z $zk_host $DASHM
$DASHA \n\n"
    # call this script again with correct args for next node
    $SOLR_TIP/bin/solr start -cloud -d node$ndx -p $next_port -z $zk_host
$DASHM $DASHA
  done

  # TODO: better (shorter) name??
```

```
  CLOUD_COLLECTION=gettingstarted

  if $noprompt ; then
    CLOUD_NUM_SHARDS=2
    CLOUD_REPFACT=2
    # if the new default config directory is available, then use it,
    # otherwise, use the legacy collection1 example
    # TODO: this will need to change when SOLR-3619 is resolved
    if [ -d "$SOLR_TIP/server/solr/configsets/schemaless" ]; then
      CLOUD_CONFIG_DIR=$SOLR_TIP/server/solr/configsets/schemaless
      CLOUD_CONFIG=schemaless
    else
      CLOUD_CONFIG_DIR=$SOLR_TIP/example/solr/collection1/conf
      CLOUD_CONFIG=default
    fi
  else
    echo -e "\nNow let's create a new collection for indexing documents in your
$CLOUD_NUM_NODES-node cluster.\n"
    read -e -p "Please provide a name for your new collection: [gettingstarted] "
USER_INPUT
    # trim whitespace out of the user input
    CLOUD_COLLECTION=`echo $USER_INPUT | tr -d ' '`

    # handle the default selection or empty input
    if [ "$CLOUD_COLLECTION" == "" ]; then
      CLOUD_COLLECTION=gettingstarted
    fi
    echo $CLOUD_COLLECTION

    USER_INPUT=
    read -e -p "How many shards would you like to split $CLOUD_COLLECTION
into? [2] " USER_INPUT
    # trim whitespace out of the user input
    CLOUD_NUM_SHARDS=`echo $USER_INPUT | tr -d ' '`

    # handle the default selection or empty input
    if [ "$CLOUD_NUM_SHARDS" == "" ]; then
      CLOUD_NUM_SHARDS=2
    fi
    echo $CLOUD_NUM_SHARDS

    USER_INPUT=
    read -e -p "How many replicas per shard would you like to create? [2] "
USER_INPUT
    # trim whitespace out of the user input
    CLOUD_REPFACT=`echo $USER_INPUT | tr -d ' '`

    # handle the default selection or empty input
```

```
    if [ "$CLOUD_REPFACT" == "" ]; then
      CLOUD_REPFACT=2
    fi
    echo $CLOUD_REPFACT

    USER_INPUT=
    read -e -p "Please choose a configuration for the $CLOUD_COLLECTION
collection, available options are: default or schemaless [default] " USER_INPUT
    # trim whitespace out of the user input
    CLOUD_CONFIG=`echo $USER_INPUT | tr -d ' '`

    # handle the default selection or empty input
    if [ "$CLOUD_CONFIG" == "" ]; then
      CLOUD_CONFIG=default
    fi
    echo $CLOUD_CONFIG

    if [ "$CLOUD_CONFIG" == "schemaless" ]; then
      if [ -d "$SOLR_TIP/server/solr/configsets/schemaless" ]; then
        CLOUD_CONFIG_DIR=$SOLR_TIP/server/solr/configsets/schemaless
      else
        CLOUD_CONFIG_DIR=$SOLR_TIP/example/example-
schemaless/solr/collection1/conf
      fi
    else
      CLOUD_CONFIG_DIR=$SOLR_TIP/example/solr/collection1/conf
    fi
  fi

  echo -e "\nDeploying default Solr configuration files to embedded ZooKeeper
using command:\n"
  echo -e "$DEFAULT_SERVER_DIR/scripts/cloud-scripts/zkcli.sh -zkhost
$zk_host -cmd upconfig -confdir $CLOUD_CONFIG_DIR -confname
$CLOUD_CONFIG\n"
  # upload the config directory to ZooKeeper
  # Extract the solr.war if it hasn't been done already (so we can access the
SolrCLI class)
  if [ ! -d "$DEFAULT_SERVER_DIR/solr-webapp/webapp" ]; then
    (mkdir -p $DEFAULT_SERVER_DIR/solr-webapp/webapp && cd
$DEFAULT_SERVER_DIR/solr-webapp/webapp && jar xf
$DEFAULT_SERVER_DIR/webapps/solr.war)
  fi
  $JAVA -Dlog4j.configuration=file:$DEFAULT_SERVER_DIR/scripts/cloud-
scripts/log4j.properties \
    -classpath "$DEFAULT_SERVER_DIR/solr-webapp/webapp/WEB-
INF/lib/*:$DEFAULT_SERVER_DIR/lib/ext/*" \
    org.apache.solr.cloud.ZkCLI -zkhost $zk_host -cmd upconfig -confdir
$CLOUD_CONFIG_DIR -confname $CLOUD_CONFIG > /dev/null 2>&1
```

```
  echo -e "Successfully deployed the $CLOUD_CONFIG_DIR configuration
directory to ZooKeeper as $CLOUD_CONFIG\n"

  # note use of ceiling logic in case of remainder

MAX_SHARDS_PER_NODE=$((($CLOUD_NUM_SHARDS*$CLOUD_REPF
ACT+$CLOUD_NUM_NODES-1)/$CLOUD_NUM_NODES))

  COLLECTIONS_API=http://localhost:$SOLR_PORT/solr/admin/collections


CLOUD_CREATE_COLLECTION_CMD="$COLLECTIONS_API?action=CR
EATE&name=$CLOUD_COLLECTION&replicationFactor=$CLOUD_REPFA
CT&numShards=$CLOUD_NUM_SHARDS&collection.configName=$CLOUD
_CONFIG&maxShardsPerNode=$MAX_SHARDS_PER_NODE&wt=json&inde
nt=2"
  echo -e "\n\nCreating new collection $CLOUD_COLLECTION with
$CLOUD_NUM_SHARDS shards and replication factor $CLOUD_REPFACT
using Collections API
command:\n\n$CLOUD_CREATE_COLLECTION_CMD\n\nFor more
information about the Collections API, please see:
https://cwiki.apache.org/confluence/display/solr/Collections+API\n"
  curl "$CLOUD_CREATE_COLLECTION_CMD"
  echo -e "\n\nSolrCloud example running, please visit
http://localhost:$SOLR_PORT/solr \n\n"
fi

exit $?
```