



wehkamp



Implementation of Zabbix at Wehkamp



Wehkamp

February, 2nd 2018

Let me introduce myself:

name: Gerhard Pegel

Age: 42

company: Wehkamp

title: Monitoring Engineer

www.wehkamp.nl



A brief history of Wehkamp

From selling mattresses to catalogues, to online, to mobile first, to...

Wehkamp is an online retailer with their main focus on the 'CEO of the Dutch Families'. Currently we have 950 colleagues working in different locations in The Netherlands.

1952 - first advert



1955 - first catalogue

1995 - first steps online

2010 - completely online

2018 - mobile first

2019 -
A shopping
experience
that makes
women
happy



Wehkamp facts

Focus on: fashion // living // beauty / baby & children - and partnerships

>400.000
different
products

>500.000
daily
visitors

670 million
Sales figures
17/18

>50%
Of all our
customers shops
using mobile

950
colleagues

>10 million
send packages

72%
Of our customers
are female

More than 2.000 brands

C&A // Vingino // Hunkemöller // Mango // Tommy Hilfiger // Scotch & Soda // ONLY
HK Living // House Doctor // Wood // Bloomingville // Zuiver // whkmp's own

Implementation of Zabbix at Wehkamp

Implementation of Zabbix at Wehkamp

Chain monitoring

"The essence of monitoring is to send the correct message with the correct content to the correct receiver about something that needs attention as well as an automated action to be able to prevent (pro-active) or to solve any kind of problems (reactive)"

This in combination with "chain monitoring" will support certain objects (infrastructure components, applications and/or databases) in relation with the company processes that will be monitored and visualised. The two most important advantages are:

1. To be able to set a better priority of a disruption
2. More transparency for the business about managing the whole chain

Implementation of Zabbix at Wehkamp

Tooling

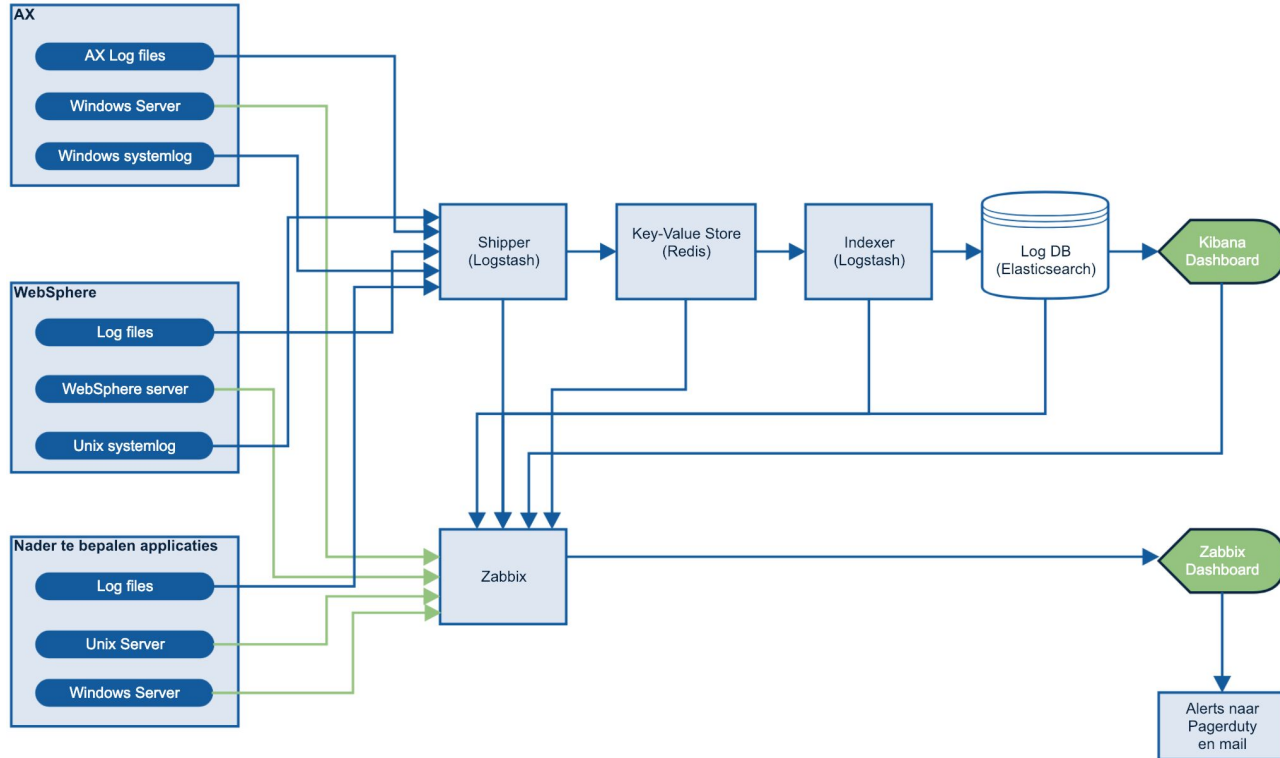
In the period of September - December 2015 an external team (**Xifeo**) and I started to setup a large part of chain monitoring for the IT infrastructure for the Lacent domain. In order to setup a central logging system (collect & visualize) we have used open source tooling like:

- Elasticsearch (search & analyze data in real time)
- Logstash (collect, enrich & transport data)
- Kibana (explore and visualize data)
- Redis (In-memory data structure store, used as database, cache and message broker)
- Grafana (was not part of the implementation but **Dennis Kanbier** told me to use it and I listened)

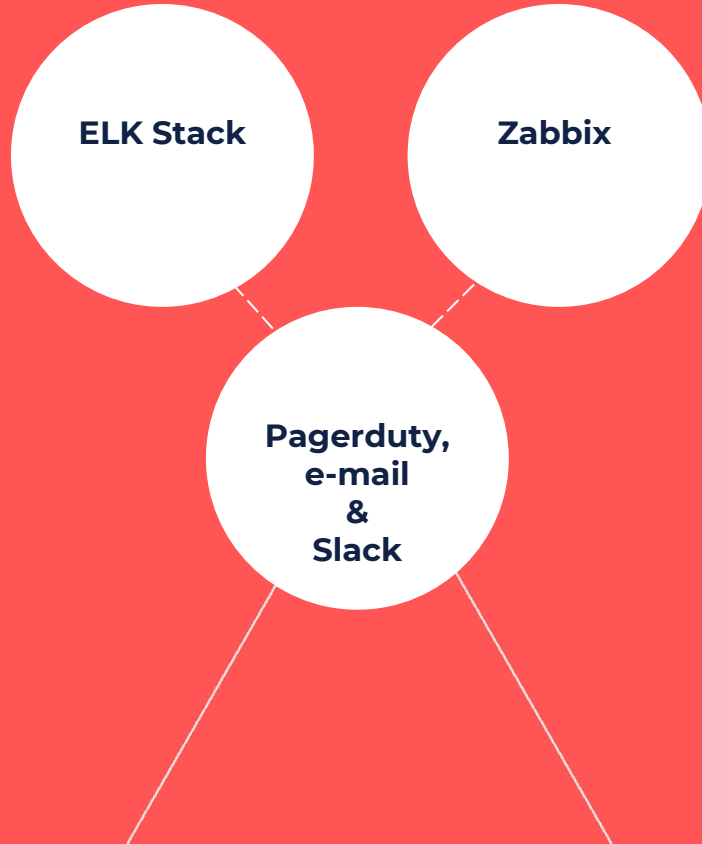
For the alerting and resource monitoring we started using Zabbix.

Implementation of Zabbix at Wehkamp

Schematic overview of our chain monitoring setup



Implementation of Zabbix at Wehkamp



Implementation of Zabbix at Wehkamp

- 1 Self healing: Elasticsearch auto cleanup
- 2 Web monitoring
- 3 ODBC monitoring
- 4 SSL Certificate monitoring
- 5 iDeal on-premise monitoring
- 6 Our future plans

Implementation of Zabbix at Wehkamp

Self healing: Elasticsearch auto cleanup

Zabbix is configured to cleanup Elasticsearch indices when disk space in /data is below 20% on any of the data nodes.

We did this while using 5 steps:

1. Monitor disk usage
2. Make a plan
3. Create the cleanup script
4. Configure a trigger action in Zabbix
5. Visualize results in Grafana

Implementation of Zabbix at Wehkamp

Self healing: Elasticsearch auto cleanup

- Monitor disk usage

Zabbix comes with the standard template "Template OS Linux". This template uses auto-discovery of mountpoints out-of-the-box. So all we need to do to actually monitor the disk usage for mount point /data is add the nodes to the "Template OS Linux" template in Zabbix.

Implementation of Zabbix at Wehkamp

Self healing: Elasticsearch auto cleanup

- Make a plan

Now each individual server gets an alert when disk space on /data < 20%. What we want to achieve is when this actually happens we clean up some old indices on the Elasticsearch cluster to free up some space. In order to this we will tell Zabbix to execute a cleanup script when free space in /data reaches < 20%.

So our plan is:

- Monitor free space in /data
- If it gets below 20% trigger an action
- The action executes a script on the ES cluster to remove the oldest indices
- Stop removing indices until a certain amount of free space is available

Implementation of Zabbix at Wehkamp

Self healing: Elasticsearch auto cleanup

- Create the cleanup script

What it does is simply remove indices until /data reaches at least 25% of free space OR it removes all indices older than 21 days.

Next we'll configure Zabbix to automatically run this script.

Please note, this script runs everyday at 23.00 via a cronjob. Just in case the cronjob fails, Zabbix will Action this script when the free disk space is going under de 20 percent.

```
1  #!/bin/bash
2  # Options
3  DATA_PARTITION="/data"
4  STOP_CLEAN_PERCENTAGE="75"
5  DID_CLEAN=false
6  MAX_KEEP=21
7
8  # Keep a maximum of $MAX_KEEP files
9  TOTAL=`curl --silent 'localhost:9200/_cat/indices?v' | egrep -o "logstash-[0-9\.]*" | wc -l`
10 while [ ${TOTAL} -gt ${MAX_KEEP} ]; do
11     OLDEST=`curl --silent 'localhost:9200/_cat/indices?v' | egrep -o "logstash-[0-9\.]*" | sort | head -1`
12     #read -p "About to delete ${OLDEST}, press any key to continue"
13     curl --silent -XDELETE "localhost:9200/${OLDEST}?pretty" >/dev/null
14     DID_CLEAN=true
15     TOTAL=$((TOTAL-1))
16 done
17
18 # Allow a maximum filesystem usage of $STOP_CLEAN_PERCENTAGE
19 USED_PERCENTAGE=`df | grep "${DATA_PARTITION}" | awk -F" " '{ print $5 }' | sed 's/%//`
20 while [ ${USED_PERCENTAGE} -gt ${STOP_CLEAN_PERCENTAGE} ]; do
21     OLDEST=`curl --silent 'localhost:9200/_cat/indices?v' | egrep -o "logstash-[0-9\.]*" | sort | head -1`
22     #read -p "About to delete ${OLDEST}, press any key to continue"
23     curl --silent -XDELETE "localhost:9200/${OLDEST}?pretty" >/dev/null
24     USED_PERCENTAGE=`df | grep "${DATA_PARTITION}" | awk -F" " '{ print $5 }' | sed 's/%//`
25     DID_CLEAN=true
26 done
27
28 # Actually purge the deleted data to free up space
29 if [ "${DID_CLEAN}" = true ]; then
30     curl --silent -XPOST 'http://localhost:9200/_optimize?only_expunge_deletes=true' >/dev/null
31
```

Implementation of Zabbix at Wehkamp

Self healing: Elasticsearch auto cleanup

- Configure a trigger action in Zabbix

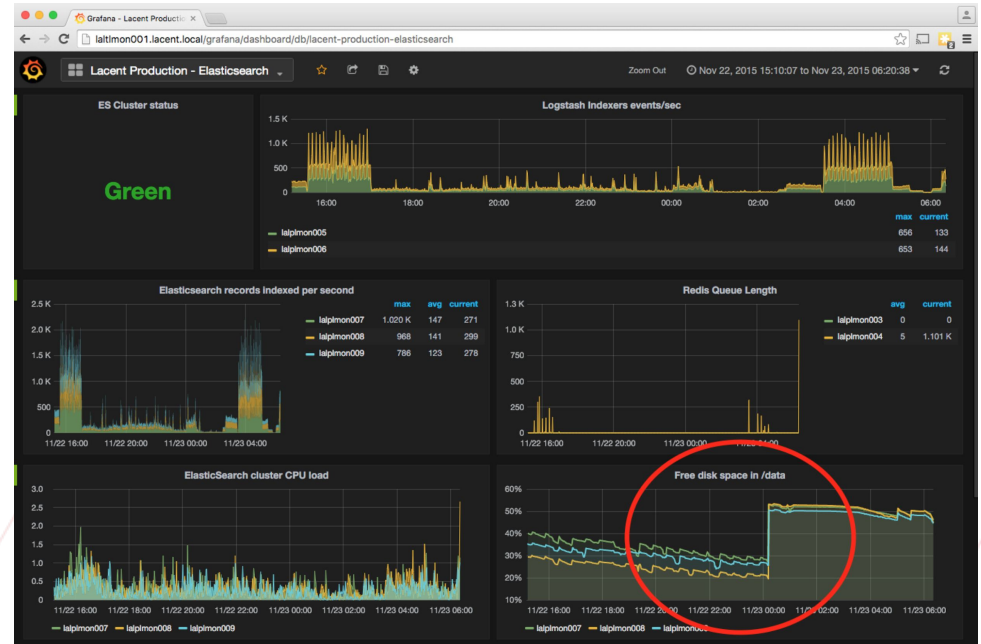
With the script in place in `/etc/zabbix/scripts/zbx_cleanup.sh` we can now tell Zabbix to execute it when the `< 20%` in data condition triggers. This is called a trigger action, as it's an action that gets executed when a trigger goes into problem state.

Implementation of Zabbix at Wehkamp

Self healing: Elasticsearch auto cleanup

- Visualize results in Grafana

With this action in place, we'll automatically clean up old indices when running low on disk space. If we look in our Grafana dashboard you can easily spot when the script is run.



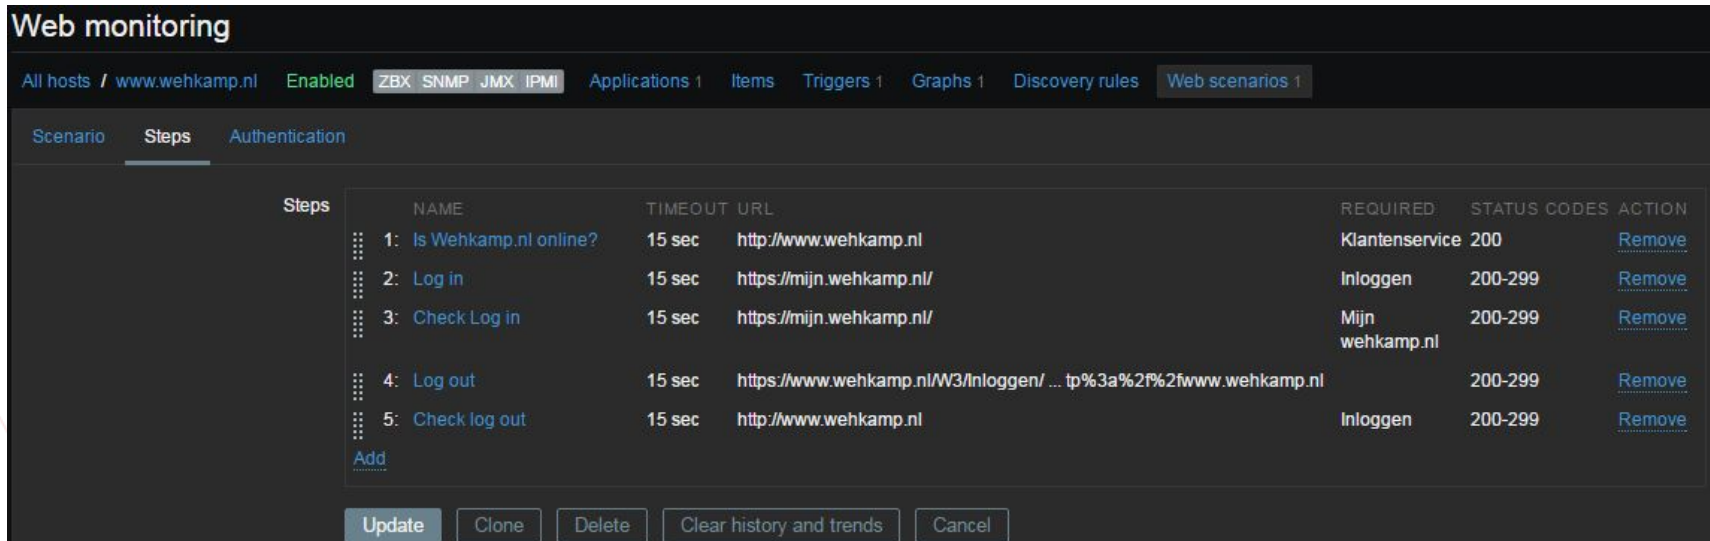
Implementation of Zabbix at Wehkamp

- 1 ~~Self healing: Elasticsearch auto cleanup~~
- 2 Web monitoring
- 3 ODBC monitoring
- 4 SSL Certificate monitoring
- 5 iDeal on-premise monitoring
- 6 Our future plans

Implementation of Zabbix at Wehkamp

Web monitoring

The scenario I've created for **www.wehkamp.nl** is more advanced. For example: the scenario called "Availability of www.wehkamp.nl" has 5 steps to check. If any of those checks fail than a trigger will be activated. Look at these steps:



The screenshot shows the Zabbix Web monitoring configuration interface for the host `www.wehkamp.nl`. The interface is in dark mode and shows the configuration for a scenario named "Availability of www.wehkamp.nl". The scenario is enabled and has 1 Web scenario. The configuration is shown in the "Steps" tab, which contains a table of 5 steps. Each step has a name, a timeout of 15 seconds, a URL, a required status code, and a "Remove" action link. The steps are: 1: Is Wehkamp.nl online? (URL: http://www.wehkamp.nl, Required: Klantenservice, Status codes: 200), 2: Log in (URL: https://mijn.wehkamp.nl/, Required: Inloggen, Status codes: 200-299), 3: Check Log in (URL: https://mijn.wehkamp.nl/, Required: Mijn wehkamp.nl, Status codes: 200-299), 4: Log out (URL: https://www.wehkamp.nl/W3/Inloggen/...tp%3a%2f%2fwww.wehkamp.nl, Required: Inloggen, Status codes: 200-299), 5: Check log out (URL: http://www.wehkamp.nl, Required: Inloggen, Status codes: 200-299). There is also an "Add" button and a "Clear history and trends" button at the bottom.

Steps	NAME	TIMEOUT	URL	REQUIRED	STATUS CODES	ACTION
1:	Is Wehkamp.nl online?	15 sec	http://www.wehkamp.nl	Klantenservice	200	Remove
2:	Log in	15 sec	https://mijn.wehkamp.nl/	Inloggen	200-299	Remove
3:	Check Log in	15 sec	https://mijn.wehkamp.nl/	Mijn wehkamp.nl	200-299	Remove
4:	Log out	15 sec	https://www.wehkamp.nl/W3/Inloggen/...tp%3a%2f%2fwww.wehkamp.nl	Inloggen	200-299	Remove
5:	Check log out	15 sec	http://www.wehkamp.nl	Inloggen	200-299	Remove

Implementation of Zabbix at Wehkamp

Web monitoring

And when everything is OK than we will see this as a result:



Implementation of Zabbix at Wehkamp

1

~~Self healing: Elasticsearch auto cleanup~~

2

~~Web monitoring~~

3

ODBC monitoring

4

SSL Certificate monitoring

5

iDeal on-premise monitoring

6

Our future plans

Implementation of Zabbix at Wehkamp

ODBC Monitoring

As you all might know, Zabbix is able to query databases but this was out of scope for the implementation of Zabbix. But... there's always a but... after the implementation we needed to start using this.

Why? Well... we had this case:

1. User enters the office before 8 am.
2. User does a query on a database before 8 am. (User needs to check this before 8 am)
3. IF result == 0 user needs to solve a problem ELSE user can have some coffee

But what happens when the result became == 0 after 8 am due to some disruption or slowness of the job processing during the night?

Implementation of Zabbix at Wehkamp

1

~~Self healing: Elasticsearch auto cleanup~~

2

~~Web monitoring~~

3

~~ODBC monitoring~~

4

SSL Certificate monitoring

5

iDeal on-premise monitoring


6

Our future plans

Implementation of Zabbix at Wehkamp

SSL Certificate monitoring

Another part of the chain monitoring was setting up monitoring about SSL certificates. By using a very small but helpful script we could visualize when SSL certificates are going to expire.



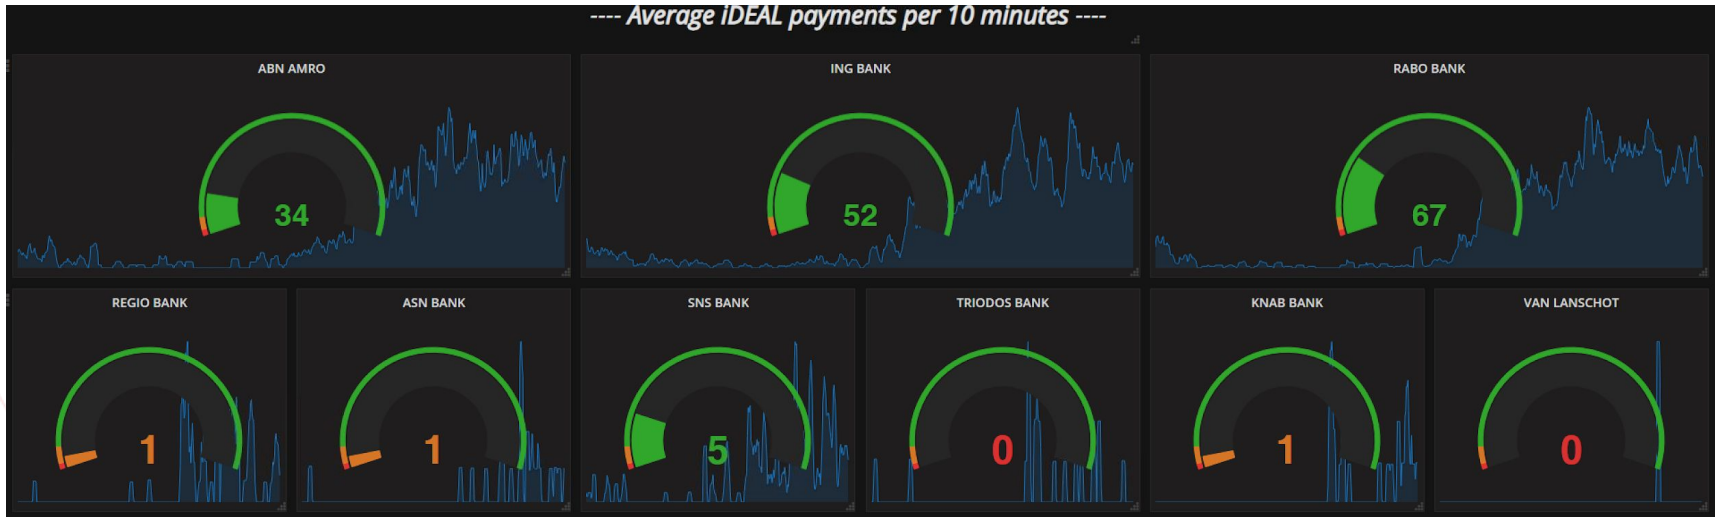
Implementation of Zabbix at Wehkamp

- 1 ~~Self healing: Elasticsearch auto cleanup~~
- 2 ~~Web monitoring~~
- 3 ~~ODBC monitoring~~
- 4 ~~SSL Certificate monitoring~~
- 5 iDeal on-premise monitoring
- 6 Our future plans

Implementation of Zabbix at Wehkamp

iDeal on-premise monitoring

Our customers have the possibility to buy their products while using iDeal as a payment option. We wanted to see if the iDeal service of our website, that's connected to a part of our on-premise infrastructure, is up and running. In case it is NOT running we created triggers and alerts that was send to our customer service team.



Implementation of Zabbix at Wehkamp

1 ~~Self healing: Elasticsearch auto cleanup~~

2 ~~Web monitoring~~

3 ~~ODBC monitoring~~

4 ~~SSL Certificate monitoring~~

5 ~~iDeal on-premise monitoring~~

6 Our future plans

**Zabbix has been implemented,
what's next?**

Our future plans...

- 1 We are going to host Zabbix in AWS
- 2 We are going to send all Zabbix history to Elasticsearch
- 3 SCOM will be replaced by Zabbix
- 4 The current bot that's connected via Slack to Zabbix will get more functionality

Thanks!

**Any
questions?**

wehkamp.nl



wehkamp



wehkamp