ZABBIX SUMMIT ONLINE /2021

BUILD ZABBIX SERVER HA CLUSTER IN 10 MINUTES

KASPARS MEDNIS

Senior Trainer, Zabbix, Latvia

01

WHY ZABBIX NEEDS HIGH AVAILABILITY

WHAT IS HIGH AVAILABILITY?

- High availability (HA) is the ability of a system to operate continuously without failing for a designated period of time
  - A single point of failure is a component that would cause the whole system to fail if it fails.
  - Redundancy enables a backup component to take over for a failed one.
  - Failures must be visible and systems have built-in automation to handle the failure on their own.

WHY ZABBIX NEEDS HIGH AVAILABILITY?

- Zabbix server without HA is a single point of failure
- If Zabbix server fails for some reason, no notifications will be sent

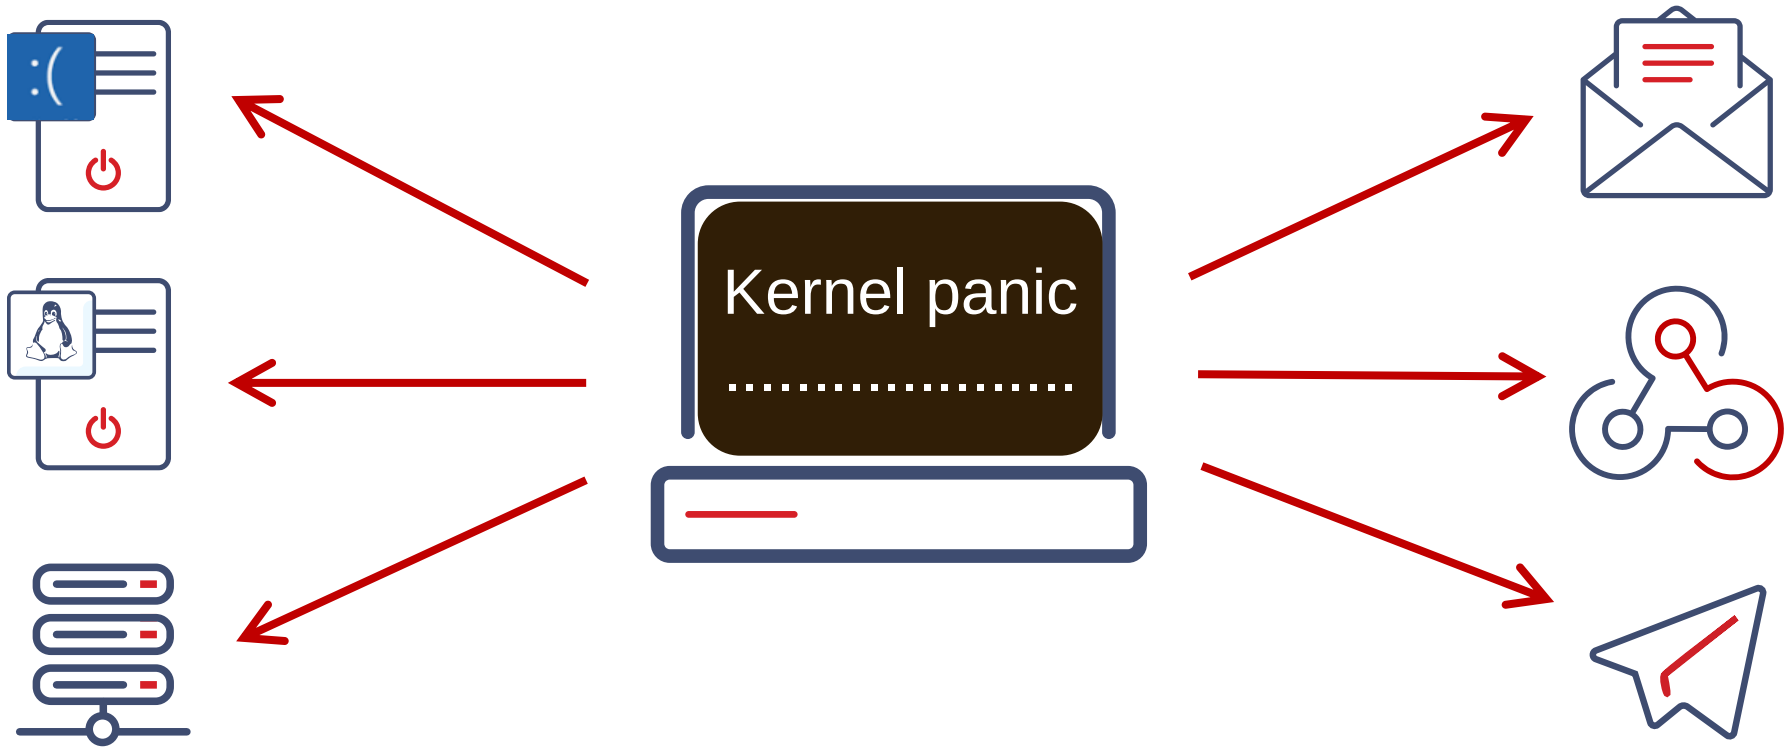

02

HOW TO CHOOSE HA FOR ZABBIX

WHICH HA SOLUTION TO CHOOSE?

- It is possible to use 3rd party HA solutions for Zabbix
- Most of them require specific knowledge for a proper set up
- Some service providers may charge additional costs for HA setup

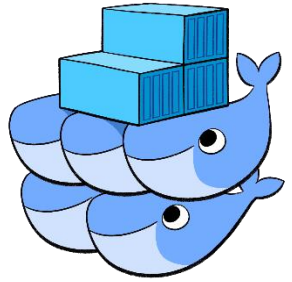

WHAT CAN GO WRONG WITH HA?

- If HA is not properly implemented failover may not happen when needed
- A split-brain scenario with two active nodes at once may be even worse
- Other unexpected scenarios like STONITH of both nodes can happen

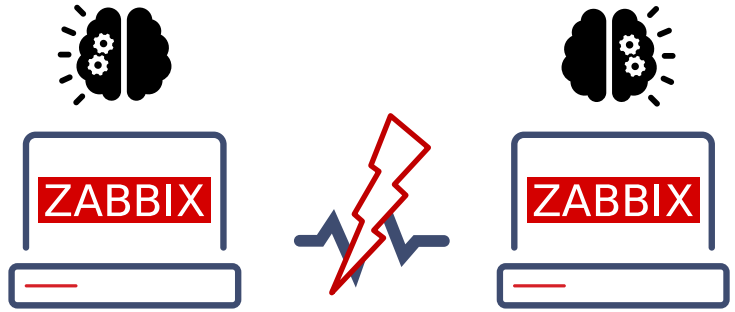

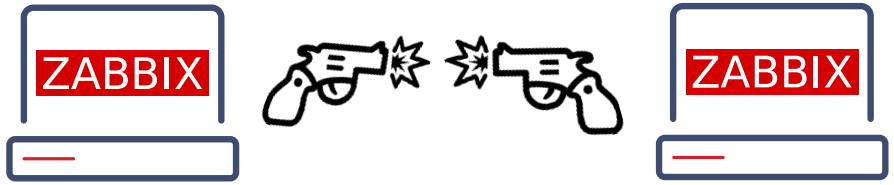

ZABBIX OFFERS NATIVE HA SOLUTION

- Easy to set up using Zabbix documentation
- Does not require expertise in HA architecture
- Officially supported by Zabbix
- Uses Zabbix database to check node status

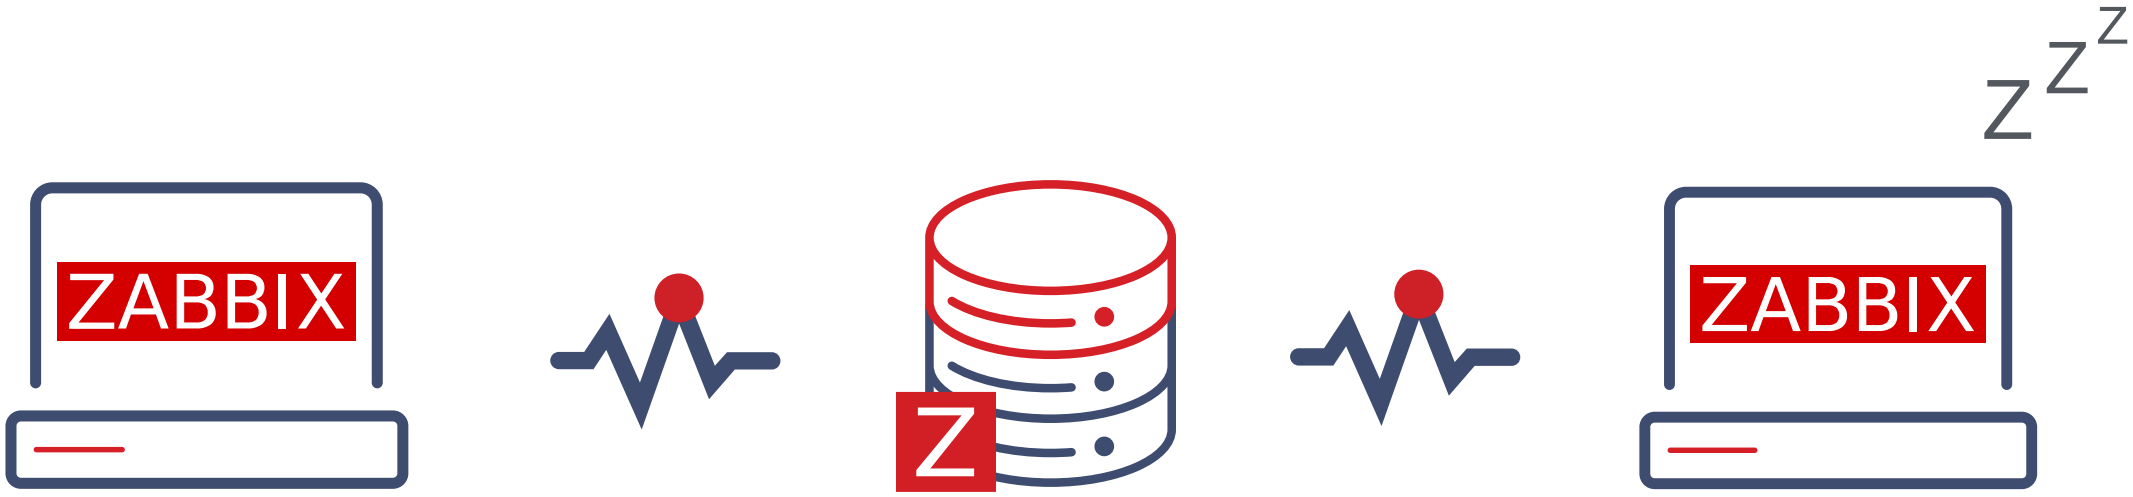

03

HOW ZABBIX CLUSTER WORKS

STARTING ZABBIX IN A CLUSTER MODE

- To start Zabbix server in HA mode, new configuration parameters were added
- Edit Zabbix configuration file (usually located at /etc/zabbix/zabbix_server.conf)
  - ✓ HANodeName to specify the name of Zabbix cluster node
  - ✓ NodeAddress to specify the address of cluster node
- Restart all Zabbix cluster nodes after making changes to configuration files

ZABBIX HA NODE NAME

- HANodeName specifies the name of the node
- Without HANodeName specified, Zabbix server will start in a standalone mode

```
## Option: HANodeName
#       The high availability cluster node name.
#       When empty server is working in standalone mode.
HANodeName=zbx-node1
```

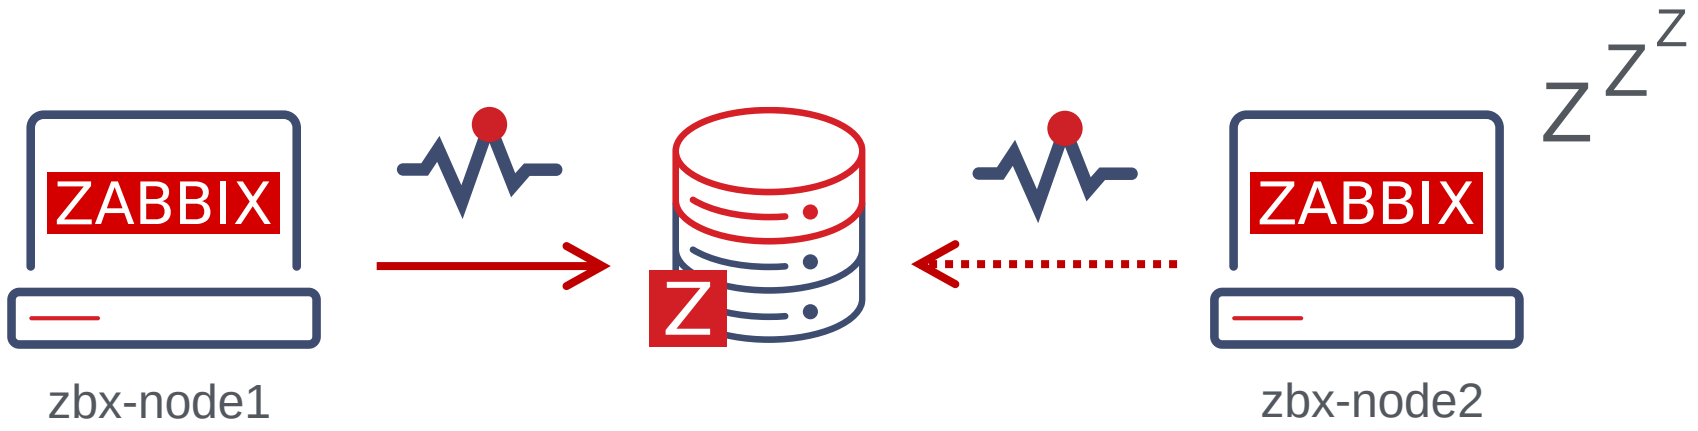

ZABBIX HA NODE ADDRESS

- Additionally, NodeAddress parameter must be specified for each node
- NodeAddress must match IP or FQDN name of Zabbix server node
- This parameter will be used by Zabbix frontend to connect to active node

```
## Option: NodeAddress
#       IP or hostname to define hoe frontend should connect to the server.
#       Format: <address>[:port]
NodeAddress=node1.example.com
```

ZABBIX FRONTEND SETUP

- Zabbix frontend will autodetect the active node
  - ✓ Frontend reads settings from the nodes table in Zabbix database
  - ✓ Node address of the active node is used as the Zabbix server address
  - ✓ Zabbix server address and port must be undefined in the frontend configuration

```
// Uncomment and set to desired values to override Zabbix hostname/IP and port.
// $ZBX_SERVER                 = '';
// $ZBX_SERVER_PORT            = '';
```

PUTTING IT ALL TOGETHER

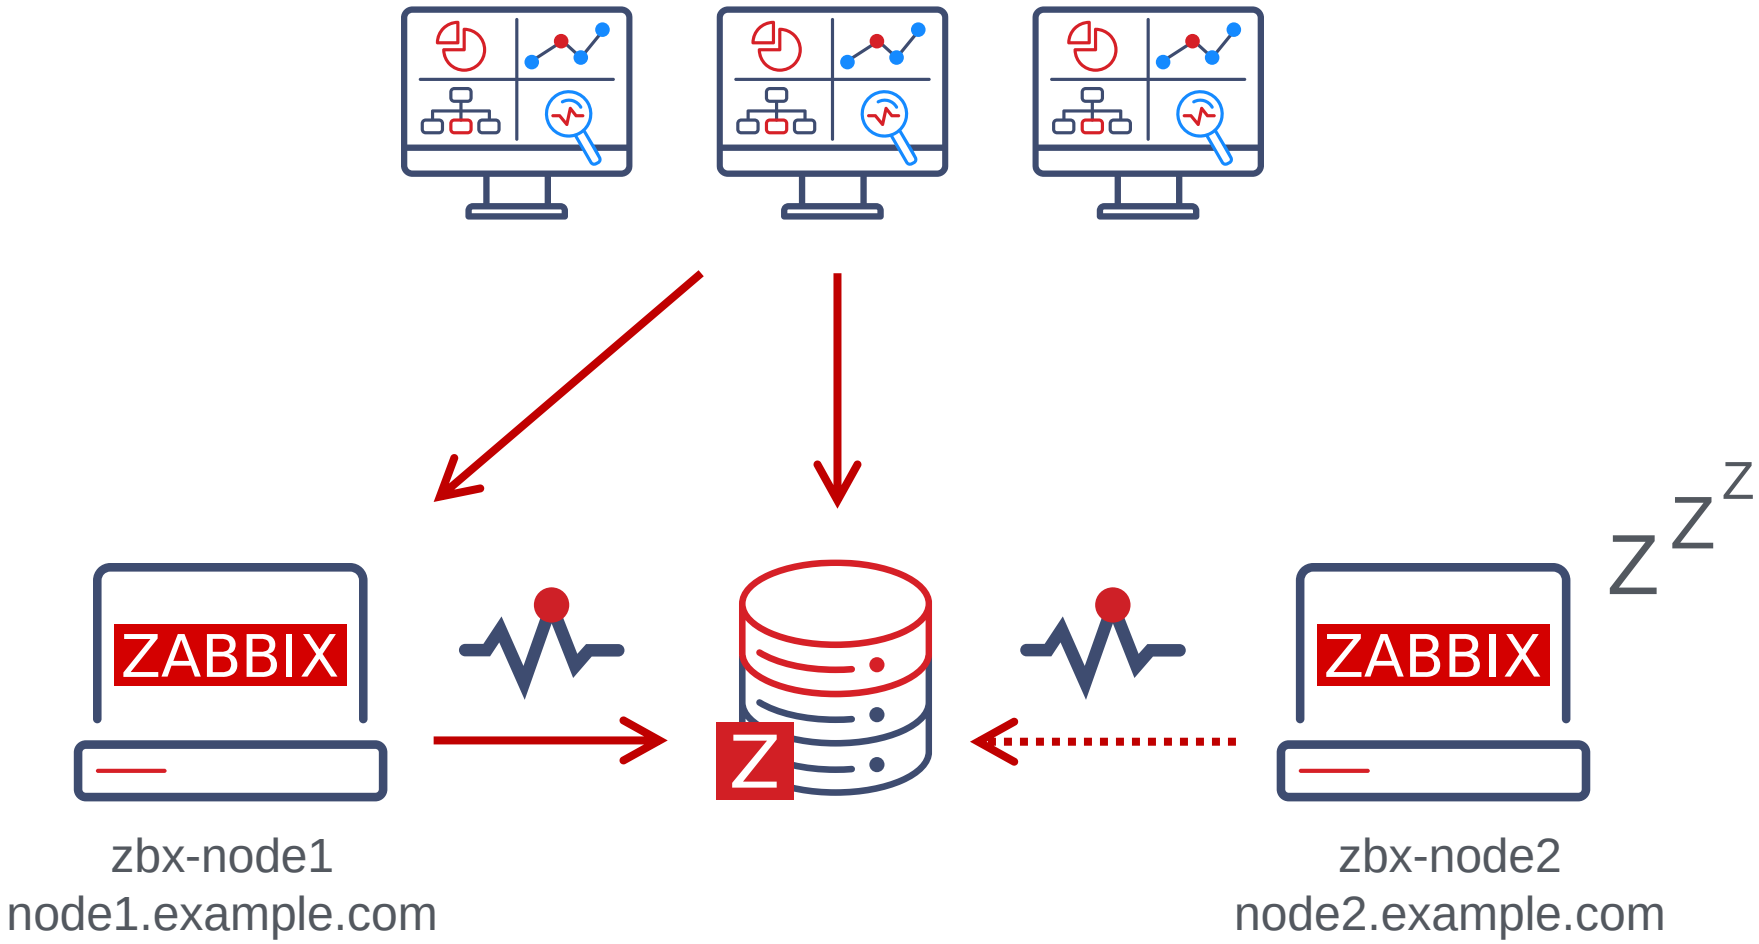

ZABBIX HA NODE TYPES

- Zabbix has 3 statuses for HA node
  - ✓ Active — only one node can be active at a time
  - ✓ Standby — multiple nodes can be in a standby mode
  - ✓ Shutdown — a node was previously detected, but is shut down now

- There is one more status for unavailable nodes
  - ✓ Unreachable — a node was previously detected, but was lost without a shutdown

ZABBIX HA MANAGER

- On every node a special process called HA Manager is started
- This is the only active Zabbix process on standby nodes started after the main process
- All other Zabbix processes will be started when failover happens

```
ps ax | grep zabbix_server
  39177 ?        S      0:00 /usr/sbin/zabbix_server -c /etc/zabbix/zabbix_server.conf
  39179 ?        S      0:00 /usr/sbin/zabbix_server: ha manager
```

ZABBIX HA NODE STATUS

- Status of all HA cluster nodes is displayed on Zabbix frontend

System information

Parameter	Value	Details
Zabbix server is running	Yes	zbx-node1:10051
Number of hosts (enabled/disabled)	1	1 / 0
Number of templates	288	
Number of items (enabled/disabled/not supported)	97	88 / 0 / 9
Number of triggers (enabled/disabled [problem/ok])	55	55 / 0 [1 / 54]
Number of users (online)	2	1
Required server performance, new values per second	1.42	
High availability cluster	Enabled	Fail-over delay: 1 minute

Name	Address	Last access	Status
zbx-node1	zbx-node1:10051	5s	Active
zbx-node3	zbx-node3:10051	3s	Stopped
zbx-node2	zbx-node2:10051	4s	Standby

SWITCHING ZABBIX HA NODE

- Zabbix will failover to another node automatically on active node stop
- There must be at least one node in standby status

```
systemctl stop zabbix-server
```

HOW A FAILOVER WORKS?

- All nodes report their status every 5 seconds
  - When a node shuts down the standby node which first detects the lost node will take over

- If the node is lost and will not respond in time
  - The clock will keep on ticking until it reaches the failover delay (1 minute by default)
  - When the failover delay is over one of the standby nodes will take over

ZABBIX CLUSTER TUNING

- It is possible to adjust failover delay using ha_set_failover_delay runtime command
  - ✓ supported range is from 10 seconds to 15 minutes

```
# zabbix_server -R ha_set_failover_delay=5m
zabbix_server: command sent successfully
```

- Nodes can be removed using ha_remove_node runtime command

```
# zabbix_server -R ha_remove_node=3
zabbix_server: command sent successfully
```

03

CONNECTING AGENTS AND PROXIES

WHAT ABOUT ZABBIX AGENTS?

- Zabbix agent requires all nodes to be written in the configuration file
- Nodes are specified in a comma-separated list

```
### Option: Server
#       List of comma delimited IP addresses or DNS names of Zabbix
#       Incoming connections will be accepted only from the hosts listed here.
Server=zbx-node01,zbx-node02
```

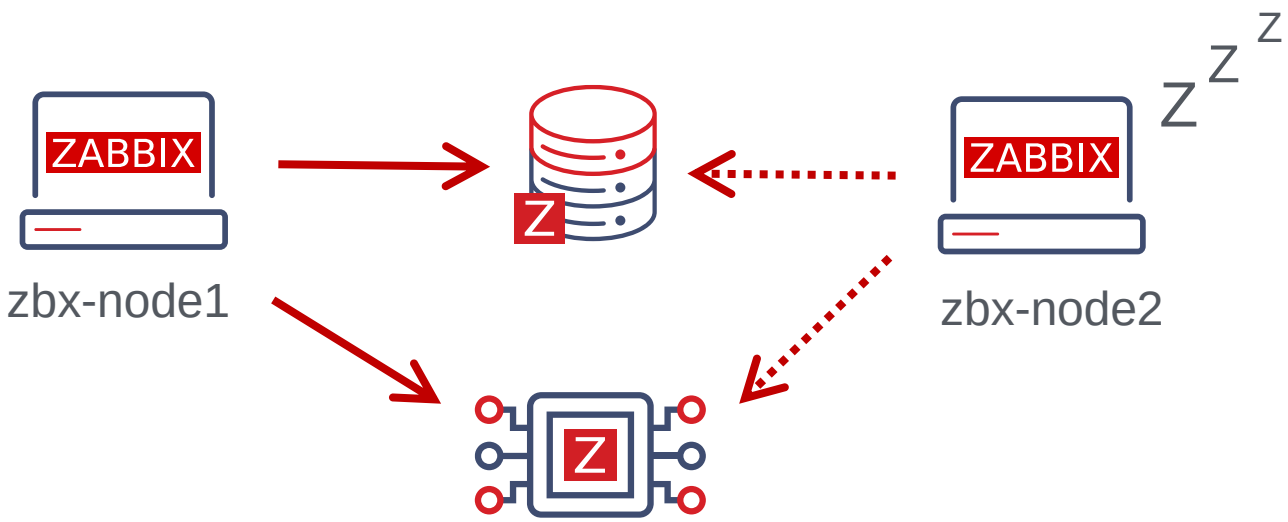

WHAT ABOUT ACTIVE CHECKS?

- Zabbix agent requires all nodes to be written in ServerActive section of the configuration file
- Cluster nodes need to be separated by a semicolon

```
### Option: ServerActive
#       List of comma delimited IP addresses or DNS names pairs of clusters of Zabbix
#       servers for active checks.
#       Cluster nodes need to be separated by semicolon.
ServerActive=zbx-node01;zbx-node02
```

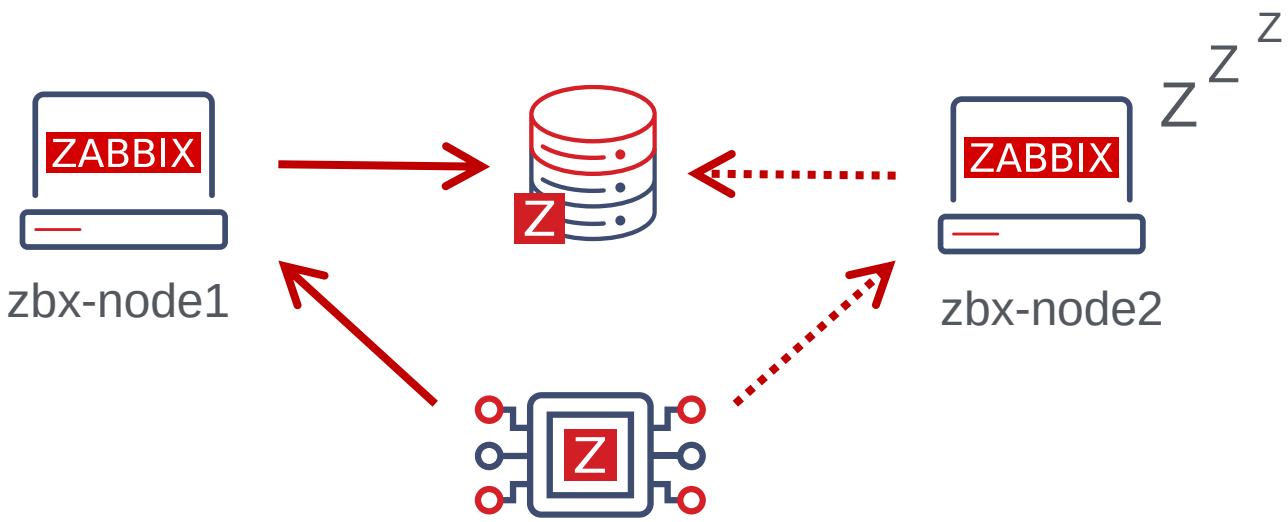

WHAT ABOUT ZABBIX PROXIES?

- Zabbix proxy in passive mode must accept connections from all Zabbix cluster nodes

```
### Option: Server
#       List of comma delimited IP addresses or DNS names of Zabbix server.
#       Incoming connections will be accepted only from the addresses listed here.
Server=zbx-node01,zbx-node02
```

- Zabbix proxy in active mode must connect to all Zabbix server cluster nodes
  - ✓ A semicolon must be used to separate node names

```
### Option: Server
#       IP address or DNS name (address:port) or cluster (address:port;address2:port) of
#       Zabbix server to get configuration data from and send data to.
#       Cluster nodes need to be separated by semicolon.
Server=zbx-node01;zbx-node02
```

SETTING UP ZABBIX HA CLUSTER

- Start Zabbix server in HA mode on all nodes (specify node name and address)
- Comment out $ZBX_SERVER in $ZBX_SERVER_PORT in frontend configuration
- Specify Server and ServerActive parameters for all Zabbix agents
- Specify Server parameter for all Zabbix proxies
- Configure other network devices for HA environment
- Enjoy !

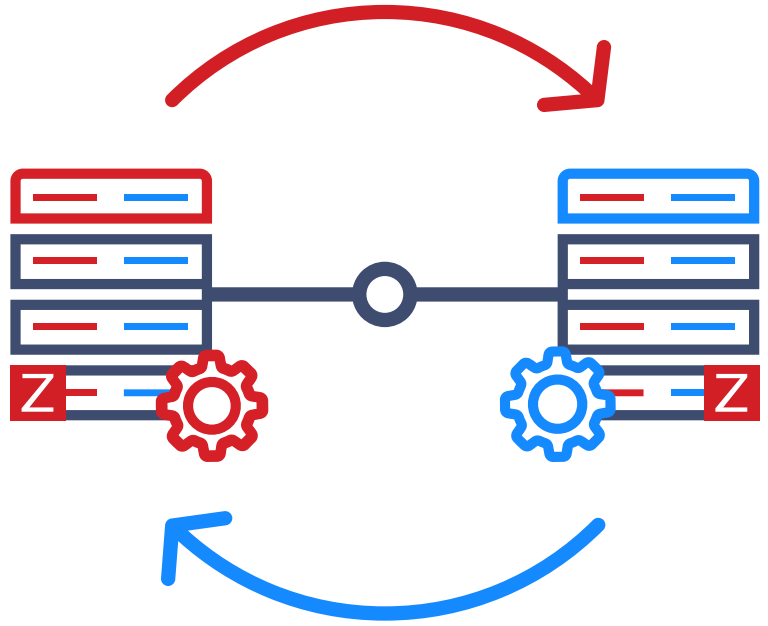

ZABBIX HA WORKSHOP AND TRAINIG

- You are welcome to participate in our Zabbix summit HA workshop
- Attend the Zabbix certified professional course to get more experience with Zabbix HA cluster


ZABBIX
SUMMIT
ONLINE /2021
Thank you!
www.zabbix.com