ZABBIX 2022 Conference BRAZIL

02 a 05 de junho de 2022 | São Paulo - SP

ZABBIX 2022 Conference BRAZIL

ZABBIX 6.0 Native High Availability

High Availability (HA) Meaning and Purpose

- Ensuring fault tolerant system
- Providing high SLA numbers
- Avoiding single point of failures
- Continuous service availability

What are available tools to achieve HA?

- Cloud service provided tools (aws, azure, etc)
- Keepalived
- PCS/Corosync/Pacemaker stack
- Database replications
- And much more…

High Availability Has Its Own Risks

- Complicated implementation
- Hard to find one template for all situations
- Maintenance
- Dependencies RHEL Example

It Really Does – Worst Real Life Scenario

- Large ZABBIX Environment
- Zabbix Engineer
- Requirement to build HA
- No previous HA experience

ZABBIX Specific Requirements

- Zabbix Server – Active/Passive only
- Zabbix Frontend – Active/Active (…)
- Zabbix Database – Depends
- How will my Agents connect to nodes?

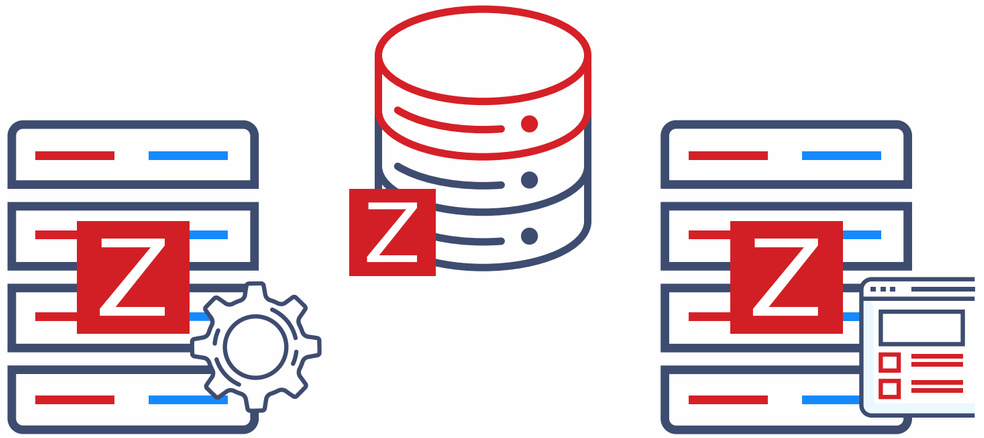

Reality

- HA Setup caused many problems
- HA Maintenance is not easy
- SLA Could be lower then without HA

ZABBIX Got You Covered

- Native Zabbix Server HA (Zabbix 6.0+)
- Simple deployment (Few minutes)
- Easy maintenance
- No dependencies
- Transparency
- Does not require expertise in HA

HOW IT ACTUALLY WORKS (Server side)

- 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 specify the name of Zabbix cluster node
- NodeAddress specify the address of cluster node

- Restart all Zabbix cluster nodes after making changes to configuration files

HOW IT ACTUALLY WORKS (Frontend side)

- Zabbix frontend will detect 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             = '';
```

HOW IT ACTUALLY LOOKS LIKE

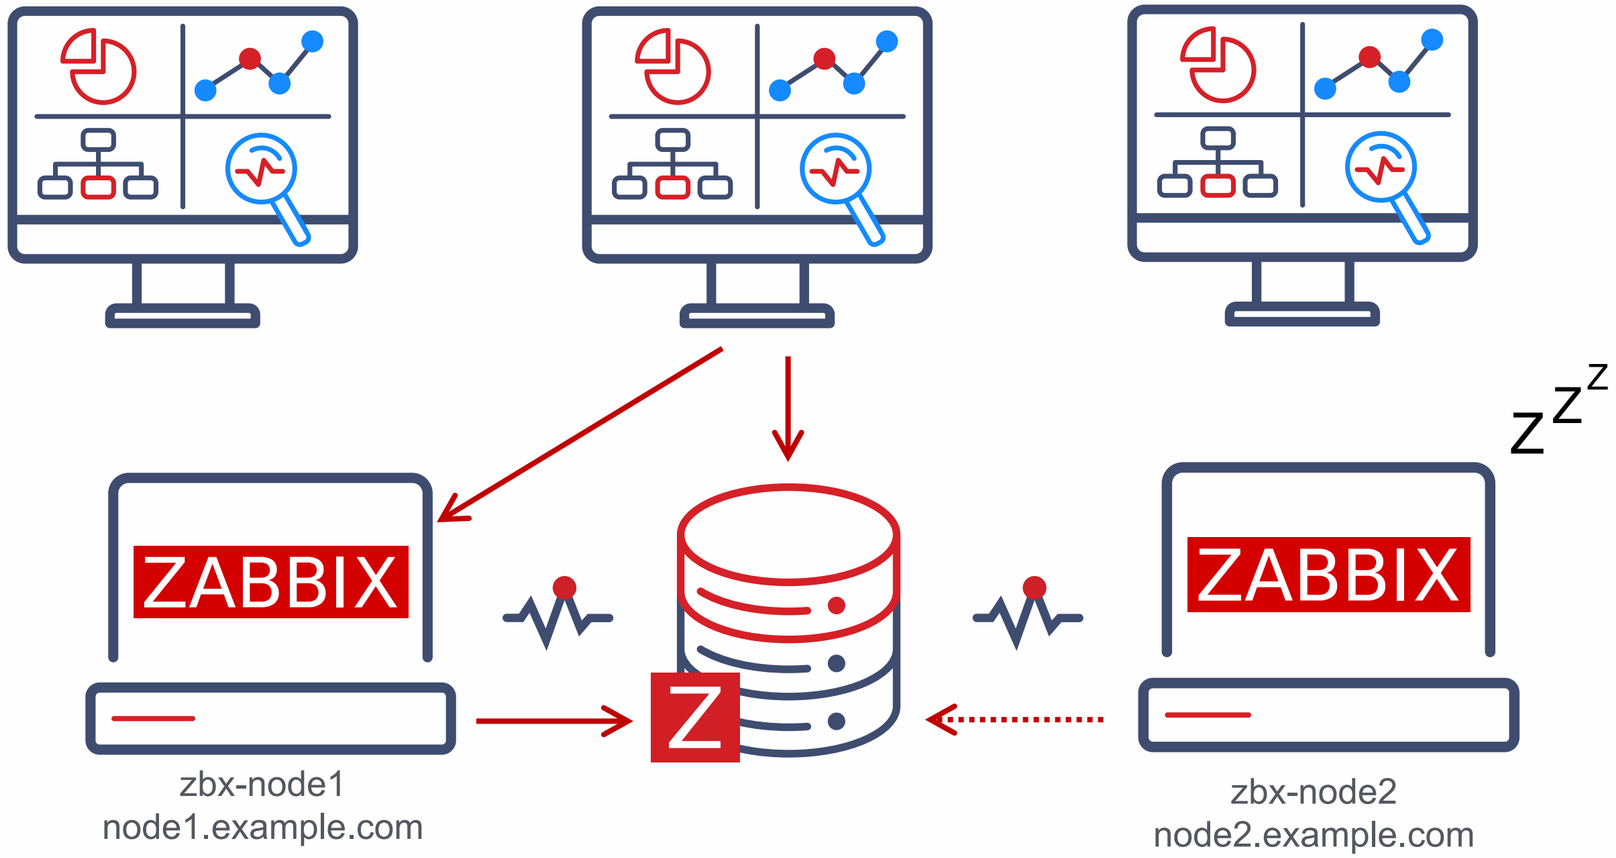

ZABBIX HA NODE TYPES

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

- Additional status for unavailable nodes
- Unreachable - a node was previously detected, but was lost without a shutdown

ZABBIX HA NODE TYPES

Status of all HA cluster nodes is displayed in 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

ZABBIX HA NODE STATUS

Status of all HA cluster nodes is displayed in 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

ZABBIX HA NODE FAILOVER

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

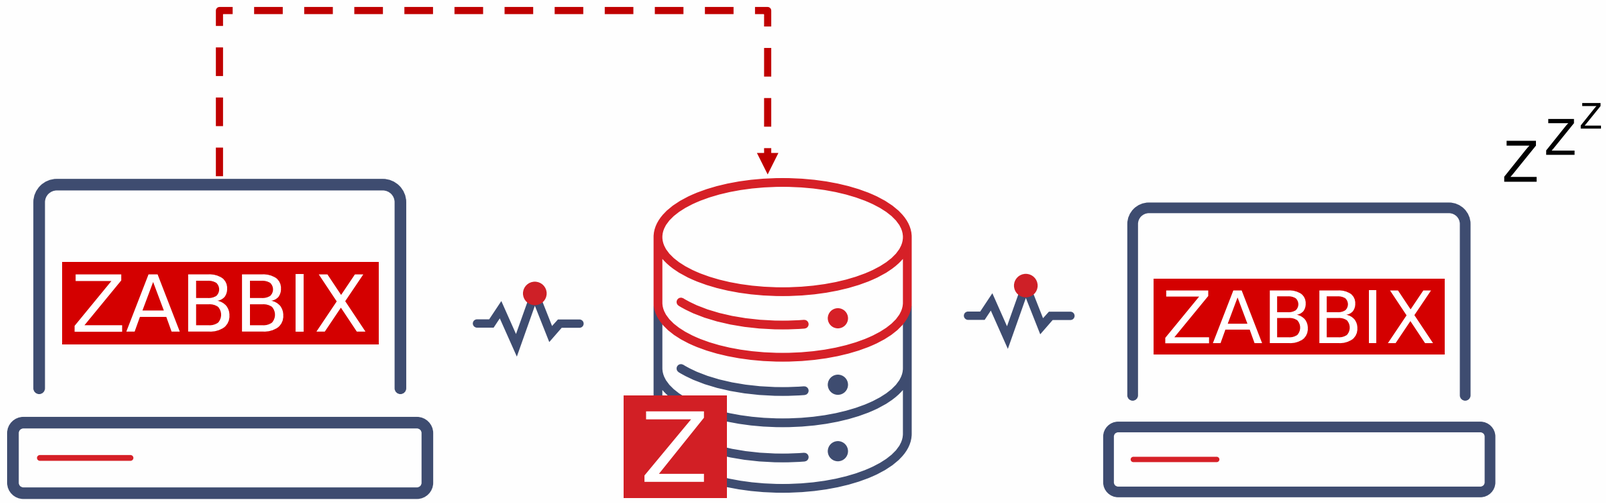

FAILOVER UNDER THE HOOD

- 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 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 HA 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 comman

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

CONNECTING ZABBIX AGENT (PASSIVE)

- 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
```

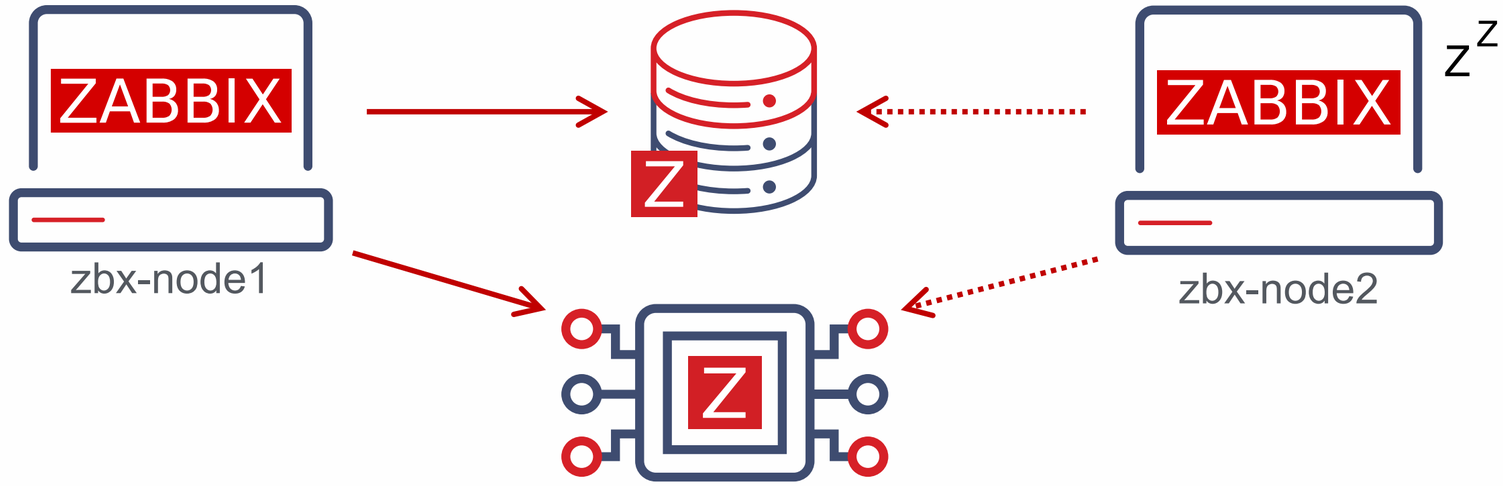

CONNECTING ZABBIX AGENT (ACTIVE)

- 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
```

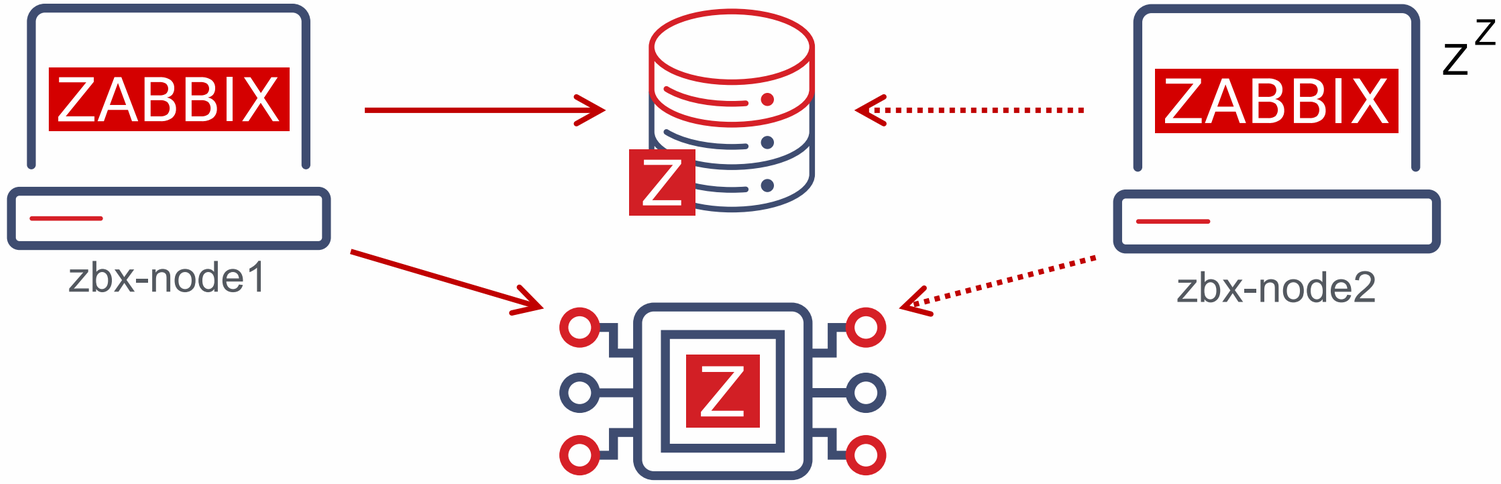

CONNECTING ZABBIX PROXY

- Zabbix proxy in passive mode must accept connection from all Zabbix cluster node

```
### 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
```

BUT WHAT ABOUT..

- ZABBIX Frontend High Availability?
- ZABBIX Database High Availability?
- Why?
- Maybe in Next releases?

ZABBIX 2022 Conference BRAZIL

THANK YOU

ZABBIX 2022 Conference BRAZIL

QUESTIONS?