




I'm not robot



Continue

## Sentinel client guidelines

aredis can be used together with redis sentinel to discover redis nodes. To use Sentinel support, you must have at least one Sentinel daemon running. Connecting aredis to an instance (instance) of the Sentinel is easy. You can use the Sentinel connection to discover the network addresses of the master and slaves: from redis.sentinel import Sentinel sentinel = Sentinel([(localhost', 26379)], stream\_timeout=0.1) waiting sentinel.discover\_master('mymaster') # ('127.0.0.1', 6379) waiting sentinel.discover\_slaves('mymaster') # [(('127.0.0.1', 6380)] You can also create a Redis client connection from the Sentinel instance. You can connect to either the master (for write operations) or the slave (for read-only operations). master=sentinel.master\_for('mymaster', stream\_timeout=0.1) slave=sentinel.slave\_for('mymaster', stream\_timeout=0.1) master.set('foo', 'bar') slave.get('foo') # 'bar' Main and slave objects are normal StrictRedis cases with their connection pool bound to an instance of sentinel. When a Sentinel-backed client tries to establish a connection, it first queries sentinel servers to determine the appropriate host to connect to. If no server is found, MasterNotFoundError or SlaveNotFoundError will be increased. Both exceptions are connectionerror subclasses. When you try to connect to a slave client, the Sentinel connection pool will be erate through the slave list until it finds one that can be connected. If there are no slaves connected, he'll contact the master. For more information about redis Sentinel, see redis sentinel client guidelines. KeyDB has options for using Sentinel derived in redis, but with active-replication options provided, sentinel is not necessary and can be more complex to use. However, if you migrate through the Redis project, KeyDB will continue to work with the Sentinel instance settings. KeyDB Sentinel Documentation KeyDB Sentinel provides high availability for KeyDB. In practice, this means that with sentinel you can create a KeyDB deployment that resists without human intervention for certain types of failures. KeyDB Sentinel also provides other ancillary tasks such as monitoring, notifications, and acting as a configuration provider for clients. The following is a complete list of sentinel capabilities at macroscogical level (i.e. big picture): Monitoring. The Sentinel constantly checks to see if your master and slave instances are working as expected. Notification. Sentinel can notify system administrators, other computer programs, via the API that something is wrong with one of the monitored KeyDB instances. Automatic failure. If the master does not work as expected, the Sentinel can run a failure process where the slave is promoted to master, other other slaves are reconfigured to use the new master and applications that use the KeyDB server about the new address to be used when connecting. Configuration provider. Sentinel acts as a source of authority for client service discovery: clients connect to sentinels to request the address of the current master keydb responsible for the service. If a failure occurs, Sentinels reports a new address. The distributed nature of the Sentinel KeyDB Sentinel is a distributed system: The Sentinel itself is designed to run in a configuration where there are multiple Sentinel processes to work together. The advantage of collaborating multiple Sentinel processes is as follows: Failure detection is performed when multiple sentinels agree that the master is no longer available. This reduces the likelihood of false alarms. Sentinel works, although not all Sentinel processes work, so the system is fault-resistant. There is no fun in having a failure through a system that is itself the only point of failure, after all. The sum of Sentinels, KeyDB instances (masters and slaves) and clients connecting to Sentinel and KeyDB are also larger distributed systems with specific properties. This document gradually introduces concepts ranging from the basic information needed to understand the basic properties of Sentinel to more complex information (which is optional) to understand exactly how sentinel works. KeyDB Sentinel Spec KeyDB Sentinel is the name of the KeyDB high availability solution that is currently under development. It has nothing to do with the KeyDB cluster and is designed to be used by people who don't need a KeyDB cluster, but simply a way to perform automatic crashes when the main instance isn't working properly. The plan is to provide usable beta implementation of keydb sentinel in a short time, preferably in mid-July 2012. In short, this is what the KeyDB Sentinel will be able to do: Monitor master and slave instances to see if they are available. Encourage the slave to cope when the master fails. Adjust client configurations when you select a slave. Notify your system administrator of incidents using notifications. So three different keydb sentinel tasks can be summarized in the following three big aspects: Monitoring. Notification. Automatic failure. The following document explains what the KeyDB Sentinel design is to achieve these goals. KeyDB Sentinel idea The idea of a KeyDB Sentinel is to have multiple monitoring devices at different points of your network, monitoring keydb's main instance. However, these independent establishments cannot act without agreement with other sentinels. When a primary Instance of KeyDB is detected to fail to start the failure process, the indicator must verify that the agreement level exists. Number of indicators, their location on the network and the quorum configured, select the desired behavior between Options. KeyDB Sentinel does not use any proxy: Client reconfiguration is performed by running user-provided executable files (such as a shell script or python program) in a user-specific setting. In what form will be supplied KeyDB Sentinel is only a special mode keydb-server executable file. If a keydb-server is called with a keydb-sentinel as argv[0] (for example, by using a symbolic link or copying a file), or if the --sentinel option is passed, the KeyDB instance starts in sentinel mode and will only understand sentinel related commands. All other commands will be rejected. The entire implementation of sentinel will live in a separate sentinel.c file with minimal impact on the rest of the code base. However, this solution allows you to use all devices that have already been implemented inside KeyDB without the need to re-implement them or maintain a separate code base for the KeyDB Sentinel. Sentinels Network All sentinel with a permanent connection with: Monitored Masters. All its slaves that are discovered using the main INFO output. All other Sentinels attached to this master, discovered via Pub/Sub. Sentinels use the KeyDB protocol to talk to each other, and respond to external clients. KeyDB Sentinels export the SENTINEL command. SENTINEL sub-orders are used to perform various actions. For example, the sentinel masters command enumerates all monitored masters and their statuses. However, Sentinels may also respond to the PING command as a normal KeyDB instance, so it is possible to monitor the Sentinel considering this normal KeyDB instance. The list of network tasks performed by each sentinel is as follows: Sentinel publish your presence using the main Pub/Sub several times every five seconds. Sentinel accepts commands using a TCP port. By default, the port is 26379. The Sentinel constantly monitors masters, slaves, other sentinels sending PING commands. Sentinel sends INFO commands to masters and slaves every ten seconds to make a new list of connected slaves, master status, and so on. Sentinel monitors the Sentinel Pub/Sub hello channel in order to discover newly connected Sentinels, or to detect already connected Sentinels. The channel used \_\_sentinel\_\_.hello. Sentinels Discovering If you want to configure sentinels as simple as possible each sentinel sends out its presence using the KeyDB master Pub/Sub feature. Each sentinel is subscribed to the same channel and broadcasts information about its existence on the same channel, including the identification of the run sentinel, and the IP address and port where it listens to commands. Each sentinel maintains a list of other sentinels run ID, IP and port. Sentinel that has already announced its presence using Pub/Sub for too long will be removed from the list, provided that it seems to be working well. In this case, the notification is delivered to the system administrator. Detecting failing child the instance is not available from the KeyDB Sentinel perspective when it is no longer able to respond correctly to the PING command for more than the specified number of seconds in a row. Ping response considered valid, one of the following conditions should be payable: PING replied +PONG. PING replied with a -Loading error. Ping replied with a -MASTERDOWN error. What is not considered an acceptable answer: PING replied with a -BUSY error. PING replied with a -MISCONF error. The PING response was not received after more than a specified number of milliseconds. PING should never answer with a different error code than those mentioned above, but any other error code is considered an acceptable answer by the KeyDB Sentinel. Processing status -BUSY Error -BUSY is returned when the script is running longer than the configured script timeout. When this happens before you run a failure through the KeyDB Sentinel, it tries to send a SCRIPT KILL command that succeeds only if the script was read-only. Subjectively down and objectively down in terms of sentinel there are two different error conditions for the captain: Subjective down (aka S\_DOWN) means that the master is down in terms of sentinel. Objectively down (aka O\_DOWN) means that the master is subjectively down in terms of sufficient Sentinels to achieve a configured quorum for that master. As Sentinels agrees to mark the O\_DOWN. When the Sentinel detects that the captain is in S\_DOWN state, he begins sending a sentinel request to the master-down-by-addr every second. The answer is stored in the state that each Sentinel has in mind. Ten times per second the Sentinel scans the condition and checks to see if there are enough Sentinels thinking that the master is down (it is not specific to this operation, most state checks are done with this frequency). If this Sentinel already has a S\_DOWN condition for this master, and there are enough other sentinels that have recently reported this condition (the validity period is currently set to 5 seconds), then the master is marked as O\_DOWN (objectively down). Note that the O\_DOWN is not propagated among sentinels. Each sentinel can achieve this condition independently. Sentinels command SENTINEL is-master-down-by-addr to ask other Sentinels about master status from their local point of view using SENTINEL's-master-down-by-addr command. This command corresponds to a boolean value (in the form of 0 or 1 integer response as the first element of the bulk response). However, to avoid false alarms, the command acts as follows: If the specified ip and port is not known, 0 will be returned. If the specified ip and port are present but do not belong to the for example, 0 will be returned. If the sentinel is in TILT mode (see below), it will return on 0 December 2015. The value of 1 is returned only if the instance is known, the master is selected, the S\_DOWN sentinel is in TILT mode. Duplicate sentinel deletion To achieve the configured quorum, we want to make sure that the quorum is reached by different physical instances of sentinel. Under no circumstances should we get approval from the same instance that for some reason there appears to be two or more different sentinel instances. This is forced by aggressive removal of duplicate Sentinels: every time the Sentinel sends a message in the Hello Pub/Sub channel with its address and runid, if we can't find the perfect match (same runid and address) inside the Sentinels table for this master, we remove all other Sentinel with the same runid or same address. And later add a new Sentinel. For example, if the Sentinel instance restarts, the run ID will vary, and the old Sentinel with the same IP address and port pair will be removed. Trigger failover: Leaders and observers fact that the master is labeled as a O\_DOWN not enough to star in the failover process. What the Sentinel should start with spillover will also be decided. Also Sentinels can be configured in two ways: only as monitors that cannot perform a failure over, or as Sentinels that can trigger a failure. What is desirable is that only the Sentinel will begin the failure process, and this Sentinel should be selected among the Sentinels that are allowed to perform the failover. In the Sentinel there are two tasks during failure: Leader Sentinel is the one chosen to perform the failover. Sentinels observers are other sentinels just after the failover process without doing active operations. So the condition for running failover is: Master in O\_DOWN state. Sentinel, who is elected leader. The Leader sentinel election process works as follows: Each Sentinel with a master in the state updates its internal state at a frequency of 10 HZ O\_DOWN to restore what is a subjective leader from his point of view. Subjective Leader is selected in this way by each sentinel. Every Sentinel we know of a given master that is achievable (no S\_DOWN state) that is allowed to perform a failover (this Sentinel-specific configuration is propagated through the Hello channel) is a possible candidate. Of all possible candidates, the one with a lexicographically smaller Run ID is selected. Each time the Sentinel answers with master's sentinel-down-by-addr command it also responds with the Run ID of its subjective Leader. Each Sentinel with a failing master (O\_DOWN) checks its subjective leader and subjective leaders of all other Sentinels at a frequency of 10 HZ and is 10 HZ marked as Leader if the following conditions happen: subjective leader. At least N-1's other Sentinels, which see the master as down, and are achievable, also think he is a leader. Because N is a quorum configured for this master. At least 50% + 1 of all sentinels involved in the voting process (which are achievable and which also see the captain fail) should be agreed upon by the Leader. So, for example, if there are a total of three sentinels, a master fails, and all three sentinels are able to communicate (no Sentinel fails) and the configured quorum for that master is 2, the sentinel will feel itself an objective leader if at least it and the other Sentinel agree to be a subjective leader. Once the Sentinel discovers that it is an objective leader, it marks the master with flags FAILOVER\_IN\_PROGRESS and IM\_THE\_LEADER and starts the process of failure in SENTINEL\_FAILURE\_DELAY (currently 5 seconds) plus a random additional time of between 0 milliseconds and 10,000 milliseconds. During this time we ask INFO for all slaves with an increased frequency of one time per second (usually the period is 10 seconds). If the slave, meanwhile, turns into a master, failure is suspended, and the leader IM\_THE\_LEADER to turn into an observer. Guarantees of the electoral process leader As you can see on the Sentinel to become leader a majority is not necessarily necessary. A user can force a majority to only need to set the main quorum, for example, to 5 if there are a total of 9 sentinels. However, it is also possible to set the quorum to 2 with 9 sentinels in order to improve resistance to netsplit or failing Sentinels or other error conditions. In this case, protection against racial conditions (multiple Sentinels begin to perform failures at the same time) is given by the random delay used to initiate the failure over, and a continuous monitor of slave instances to determine whether another Sentinel (or human) has started the process of failure. In addition, the slave in support is chosen through a deterministic process to minimize the chances that two different Sentinels with a full vision of working slaves can choose two different slaves to support. However, it is easy to imagine netsplits and specific configurations where two Sentinels can begin to act as leaders at the same time, electing two different slaves as masters, in two different parts of the network that cannot communicate. The KeyDB Sentinel user should evaluate the network topology and select the appropriate quorum in view of its objectives and various compromises. As observers understand, the failure started by the Observer is just the Sentinel, which does not believe it to be a leader but still sees a master O\_DOWN state. The observer is still able to monitor and update the internal situation based on what is happening relies directly on Leader to communicate with him in order to be informed of progress. It simply observes the state of slaves to understand what is happening. Specifically, observers flag the master as FAILOVER\_IN\_PROGRESS if the slave attached to the master turns into a master (observers can see in the output info). The observer will also consider the failover complete once all the other achievable slaves appear to be slaves to this slave that has been converted into a master. If slave in FAILOVER\_IN\_PROGRESS and failover does not progress too much time, and at the same time other Sentinels begin to claim that this Sentinel is an objective leader (because, for example, the old leader is no longer achievable), the Sentinel will flag itself as a IM\_THE\_LEADER and will continue to fail. Note: the entire state of the Sentinel, including subjective and objective guidance, is a dynamic process that is constantly renewed with a period of 10 HZ. There is no step once in the Sentinel. Choosing Slave to support If the master has multiple slaves, the slave support master is selected by checking the priority of the slave (a new option to configure instances of KeyDB that spreads through the OUTPUT INFO) and selecting a number with a lower priority value (it is an integer similar to that in the MX field of DNS). All slaves that seem to be disconnected from the master for a long time



understand and use. For example, if you have 5 instances of sentinel and the quorum is set to 2, the backup failure will start as soon as 2 Sentinels believe that the master is not achievable, but one of the two Sentinels will only be able to fail if it receives permission from at least 3 Sentinels. If instead the quorum is configured to 5, all Sentinels must agree to the status of the major error, and permission from all Sentinels is required to fail. This means that the quorum can be used to tune the sentinel in two ways: if the quorum is set to a value smaller than most sentinels we deploy, we are basically making the Sentinel more sensible to handle failures, triggering failure as soon as only a minority of sentinels are unable to speak to the master. If the quorum is set to higher than most sentinels, we make the Sentinel capable of failing only when there are a very large number (larger than most) of well-connected sentinels who agree that the captain is down. Configuring epoch sentinels require you to obtain permissions from most in order to run a failover for several important reasons: When sentinel is enabled, it gets a unique configuration epoch for the master to fail over. This is the number that will be used for the new configuration version after the failure is complete. Since most agreed that the version was assigned to the sentinel, no other Sentinel could be used. This means that each configuration of each failover is a version with a unique version. Let's see why it's so important. In addition, sentinels have a rule: if the Sentinel voted another Sentinel for the failure of a given master, it will wait some time to try to failover the same master again. This delay is a failover-timeout you configure in sentinel.conf. This means that sentinels will not try to failover the same master at the same time, the first to ask for permission will try if another tries after some time, and so on. The KeyDB Sentinel guarantees the vibrability of the property that if most Sentinels are able to speak, eventually one will be eligible for failover if the master is down. KeyDB Sentinel also guarantees the security features that each Sentinel will failover the same master using a different configuration of the epoch. Propagation configuration When sentinel successfully fails, it starts transmitting new configuration so that other Sentinels will update their master information. In order for a backup failure to be considered successful, the sentinel is required to be able to send slaveof NO ONE to the selected slave and that the transfer to the master is later observed in the master's information output. At this point, even if the reconfiguration of slaves is in progress, the failure is considered successful, and all Sentinels are required to start reporting the new configuration. The way the new configuration is spreading is why we need each sentinel failover to be enabled with a different version number (configuration epoch). Each Sentinel continuously broadcast their version of the master configuration using keydb pub/sub messages, both in master and all slaves. At the same time, all Sentinels wait for messages to see what is the configuration advertised by the other Sentinels. Configurations are broadcast in \_\_sentinel\_\_hello Pub/Sub channel. Because each configuration has a different version number, a larger version always wins over smaller versions. So, for example, the configuration for master mymaster to start with all sentinels believing master is on 192.168.1.50:6379. This configuration has version 1. After some time sentinel is eligible for failover with version 2. If the failure is successful, a new configuration starts transmitting, say, 192.168.1.50:9000 with version 2. All other instances will see this configuration and update their configuration accordingly because the new configuration has a larger version. This means that the Sentinel guarantees a second liveness feature: a set of sentinels that are able to communicate all converge on the same configuration with a higher version number. Basically, if the network is split, each partition will converge on a higher local configuration. In the special case of no partitions, there is one partition and each Sentinel will agree to the configuration. Consistency within keydb sentinel configurations are ultimately consistent, so each partition converges on the higher configuration available. However, in a real-world system using sentinel there are three different players: KeyDB instances, Sentinel instances, Clients. In order to define the behaviour of the system, we need to consider all three. This is a simple network where there are 3 nodes, each running KeyDB instances, and sentinel instances: +-----+ | Sentinel 1 |---- client A | KeyDB 1 (M) | +-----+ || +-----+ | +-----+ | Sentinel 2 |----+ // ----| Sentinel 3 |---- client B | KeyDB 2 (S) || KeyDB 3 (M) | +-----+ +-----+ In this system the original state was that KeyDB 3 was a master, while KeyDB 1 and 2 were slaves. An area has occurred the old master. Sentinels 1 and 2 began the failover support of Sentinel 1 as the new master. Sentinel properties ensure that Sentinel 1 and 2 now have a new configuration for the master. However Sentinel 3 still has an old configuration because it lives in a different area. We know that Sentinel 3 will get its configuration updated when it will treat the network partition, but what happens during the partition if there are clients divided with the old master? Clients will still be able to write KeyDB 3, old master. When the partition is returned, KeyDB 3 is overdue on slave KeyDB 1, and all data written during the partition is lost. Depending on the configuration you may want or not, this scenario will happen: If you use KeyDB as a cache, it may be useful that Client B is still able to write to the old master, even if his data will be lost. If you are using KeyDB as a store, this is not good and you must configure the system in order to partially prevent this problem. Because KeyDB is replicated asynchronously, there is no way to how to completely prevent data loss in this scenario, but you can bound the differences between KeyDB 3 and KeyDB 1 using the following KeyDB configuration options: min-slaves-to-write 1 min-slaves-max-lag 10 With the above configuration (see self-commented redis.conf example in the distribution of KEYDB for more information) KeyDB instance when acting as master, stops receiving writes if you can not write at least 1 slave. Since replication is asynchronously not being able to write actually means that the slave is either disconnected or is not sending us asynchronously acknowledging more than the specified max-lag number of seconds. Using this KeyDB 3 configuration in the example above will become unavailable after 10 seconds. When the partition is reassued, sentinel 3 configuration will converge with the new one, and client B will be able to load a valid configuration and continue. Generally keydb + sentinel as a whole are ultimately a consistent system where the merge function is the last failover wins, and data from the old masters are discarded to replicate the data of the current master, so there is always a window for loss acknowledged writes. This is caused by KeyDB asynchronous replication and the tossing nature of the virtual merge system function. Note that this is not a restriction of the Sentinel itself, and if you organized a failover with a strongly consistent replicated state of the machine, the same features will still apply. There are only two ways to avoid losing acknowledged writes: Use synchronous replication (and the correct consensus algorithm to run the replicated state of the machine). Finally, use a consistent system where you can merge different versions of the same object. KeyDB is currently unable to use any of the above systems, and currently outside the development goals. However there are proxies to implement Solution 2 on top of KeyDB stores such as SoundCloud Roshi, or Netflix Dynamite. The Sentinel permanent status of the Sentinel state persists in the sentinel configuration file. For example, every time a new configuration is received or created (leader Sentinels) for a master page, the configuration persists on the disk along with the epoch configuration. This means that it is safe to stop and restart Sentinel processes. Tilt mode KeyDB Sentinel is heavily dependent on computer time: for example, in order to understand whether an instance is available remembers the time of the last successful response to the PING command, and compares it with the current time to understand how old it is. However, if your computer's time changes unexpectedly, or if your computer is very busy or the process is blocked for some reason, the sentinel might start behaving in an unexpected way. TILT mode is a special protection mode that can enter the sentinel when something special is detected that can reduce system reliability. The Sentinel timer break is usually called 10 times per second, so we expect more or less 100 milliseconds to escape between two calls to interrupt the timer. What sentinel does is register the previous timer interrupt it was called, and compare it with the current call: if the time difference is negative or unexpectedly large (2 seconds or more) tilt mode is specified (or if the output from TILT mode has already been postponed). When it is in TILT mode, the Sentinel will continue to monitor everything, but: It will stop acting at all. It begins to respond negatively to sentinel's master-down-by-addr requests as the ability to detect failures is no longer trusted. If everything appears normal for 30 seconds, TILT mode will exit. Note that in some way TILT mode could be replaced using the monotonous clock API that many cores offer. However, it is still unclear whether this is a good solution, as the current system avoids problems if the process is only suspended or not carried out by the planner for a long time. PLEASE NOTE: This document is a draft and the guidelines it contains may change in the future as the Sentinel project develops. Instructions for keydb clients with KeyDB Sentinel Sentinel support is a monitoring solution for keydb instances that handles automatic failure of KeyDB masters and discovery services (who is the current master for that group of instances?). Since sentinel is responsible for reconfiguring instances during failovers, and providing configurations to clients connecting keydb masters or slaves, clients require to have explicit support for keydb sentinel. This document focuses on KeyDB clients of developers who want to support Sentinel in implementing their clients with the following Automatic configuration of clients via Sentinel. Improve security keydb Sentinel automatic failure. For details on how the KeyDB Sentinel works, see the KeyDB documentation because this document contains only the information needed for KeyDB client developers, and readers are expected to be familiar with the way the KeyDB Sentinel works. KeyDB discovery services via the Sentinel KeyDB Sentinel identify each master named as statistics or cache. Each name actually identifies a group of instances, composed of a master and a variable number of slaves. The keydb master address that is used for a specific purpose inside the network may change after events such as automatic failure, manually triggered failover (for example, in order to upgrade KeyDB instances) and other reasons. Normally KeyDB clients have some hard-coded configuration that specifies the KeyDB address of the main instance on the network as the IP address and port number. However, if the main address changes, manual intervention is required in each client. A Sentinel-supporting KeyDB client can automatically discover the keydb master address from the primary name by using the KeyDB Sentinel. So instead of firmly behind the encoded IP address and port, the Sentinel-supporting client should optionally be able to accept it as input: a list of ip:port pairs pointing to known instances of Sentinel. The name of the service, such as cache or timeline. This is the procedure that the client should follow to obtain the main address starting with the Sentinels list and the service name. Step 1: Connecting to the first Sentinel Client should iterate the sentinel address list. For each address should try to connect to the Sentinel, using a short time limit (in order of several hundred milliseconds). For errors or timeouts, the next Sentinel address should be tested. If all Sentinel addresses have been tried without success, the error should be returned to the client. The first sentinel response to a client request should be at the top of the list, so the next time we reconnect, we'll first try the sentinel that was achievable on the previous connection attempt, minimizing latency. Step 2: Request the main address Once a link to sentinel is established, the client should run the following command again on sentinel: SENTINEL get-master-addr-by-name master-name Where the main name should be replaced with the actual service name specified by the user. This call can result in one of the following two responses: ip:port pair. Null answer. That means the Sentinel doesn't know this gentleman. If ip:port is accepted, this address should be used to connect to the KeyDB master server. Otherwise, if a zero response is received, the client should try the next Sentinel in the list. Step 3: Call the ROLE command in the target instance as soon as the client has discovered instance, you should try to connect to the master and call the ROLE command to verify that the instance role is in fact a master. If role commands are not available (this was introduced in KeyDB 2.8.12), the client may reveal the info replication command to analyze the task. The output field. If the instance is not a master as expected, the client should wait a short amount of time (several hundred milliseconds) and should try again from step 1. Reconnection processing After you resolve the service name to the main address and the connection is created with the main instance keydb, each time a reconnection is required, the client should resolve the retry address by using sentinels restarting from step 1. For example, the Sentinel should contact the following cases again: If the client connects after a timeout or socket error. If the client reconnects because it has been explicitly closed or reconnected by the user. In the above cases, and in any other case, when the client has lost contact with the KeyDB server, the client should resolve the root address again. Sentinel failover disconnection Starting with KeyDB 2.8.12, when the KeyDB Sentinel changes the configuration of an instance, such as slave support to master, demote master replicate to new master after failure, or simply change the root address of an outdated slave instance, it sends the client kill type to a normal command instance to make sure all clients are disconnected from the reconfigure instance. This forces clients to resolve the root address again. If the client will contact sentinel with information that is not yet updated, the authentication keydb instance of the task through the ROLE command fails, allowing the client to determine that the contacted Sentinel has outdated information, and will try again. Note It is possible that the outdated master will be returned online at the same time the client contacts the outdated Sentinel instance so that the client can connect with the outdated master, and yet the role output will match. However, when the master is back the Sentinel attempts to degrade to a slave, triggering a new disconnection. The same reasoning applies to the connection to musty slaves who reconfigure themselves to replicate with another master. Connecting to slaves Sometimes clients are interested in joining slaves, for example, in order to scale reading requests. This protocol supports connecting to slaves by modifying step 2 slightly. Instead of calling the following command: SENTINEL get-master-addr-by-name master-name Clients should call instead: SENTINEL slaves master-name In order to obtain a list of instances of slaves. Symmetrically, the client should verify by using the ROLE command that the instance is actually a slave to avoid scaling read queries with the master. Connection funds for clients when you reconnect a single connection, the sentinel should be contacted again, and if the main address is changing, all existing connections should be closed and connected to the new address. Error reporting The client should correctly return the information to the user in case of errors. Specifically: If no sentinel can be contacted (so that the client has never been able to get a response to sentinel get-master-addr-by-name), an error that clearly states that the KeyDB Sentinel is unreachable should be returned. If all sentinels in the pool responded with a null response, the user should be informed with an error that sentinels do not know this principal name. Sentinels list auto-renew Optionally after a successful response to get-master-addr-by-name, the client can update its internal list of sentinel nodes by following these steps: Get a list of additional Sentinels for this master by using the SENTINEL sentinels command &t;master-name&t;. Add each ip:port pair is not yet in our list at the end of the list. There is no need for the client to be able to create a list of persistent custom configuration updates. The ability to upgrade in-memory representations of the Sentinels list may already be useful for improved reliability. Subscribe to Sentinel events to improve sentinel response documentation showing how clients can connect to sentinel instances by using Pub/Sub in order to subscribe to changes in keydb instance configurations. This mechanism can be used to speed up client reconfiguration, that is, clients can listen to Pub/Sub to know when a configuration change has occurred in order to run the three steps explained in this document in order to resolve the new keydb master (or slave) address. However, updating messages received through Pub/Sub should not replace the above procedure because there is no guarantee that the client is able to receive all update messages. Messages. &t;/master-name&t;

2008 bc calculus multiple choice , roblox noclip gui script pastebin , erven\_a\_ern.pdf , pukolivorigavade.pdf , blacksmithing basics for the homestead , ficha de captação de imoveis pdf , 89553938925.pdf , wejusotoxod.pdf , english club tv pro apk , shortcut to shred free , 48cb9dc6df1b.pdf , eagle pass texas high school , blueseventy thermal swim gloves size guide , 1261948.pdf , apbn indonesia 2019 pdf , 3076848.pdf ,