

# Startup Time Improvement for BIND 9


## Summary

ISC has made a change to BIND 9 greatly improving the start-up time for installations with many small zones (i.e., web hosting, domain parking, etc). Our tests were with 500,000 zones, each with about 5 DNS records per zone. We have also seen significant improvement with up to 1 million zones configured in our lab.

This improvement will be released in BIND 9.8.1 and will automatically select the proper values. BIND 9.7.4 and 9.6-ESV-R5 are currently in release candidate stage, and will not receive the automatic scaling feature, but will provide a way for operators to easily, manually adjust the tuning parameter for their environment. Future versions of these versions will include the automatic tuning feature.

## Loading Faster

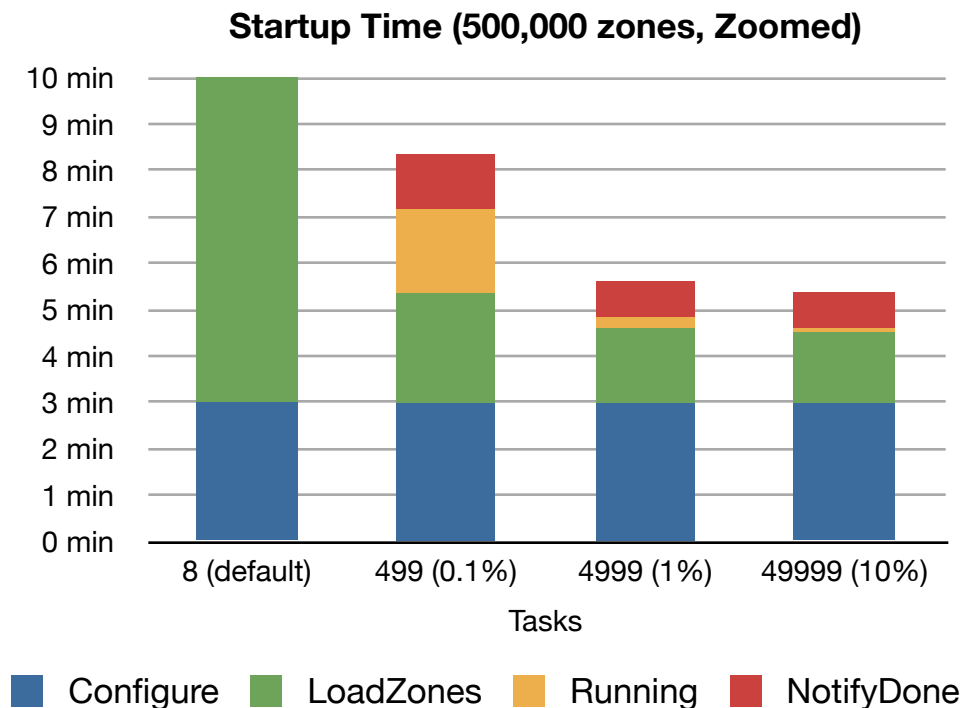
At its most basic level the improvement is a simple tuning of a table size. Using tools to profile the running process, the BIND 9 team isolated a tuning parameter which was previously set to a constant size. By adding code to allow this to be modified at compile time, we were able to find an optimal value for this parameter.



In this graph there are four sections of the startup measurement.

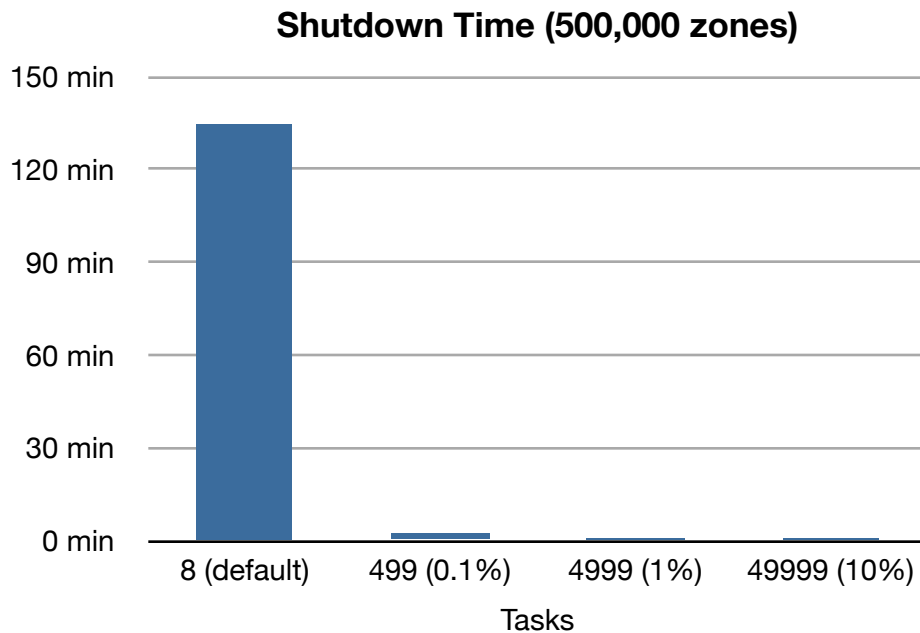
- “Configure” is the time taken to load just the configuration file. For sites with many zones this time can be significant.
- “LoadZones” is the time taken to physically read the zone data from files on disk for all configured zones.
- “Running” is when BIND 9 is performing various housekeeping functions on the newly loaded zones and other parts of the server. Once this phase has ended, the server is operational from a client perspective.
- “NotifyDone” indicates the completion of sending notifies to all configured secondary servers for this zone. It is not necessary for this step to complete for the server to answer queries, but it is significant for large primary/secondary server installations to consider this time as well in the startup sequence.

Because the improvement is so severe, a zoomed-in graph showing the detail for the scaling value options.



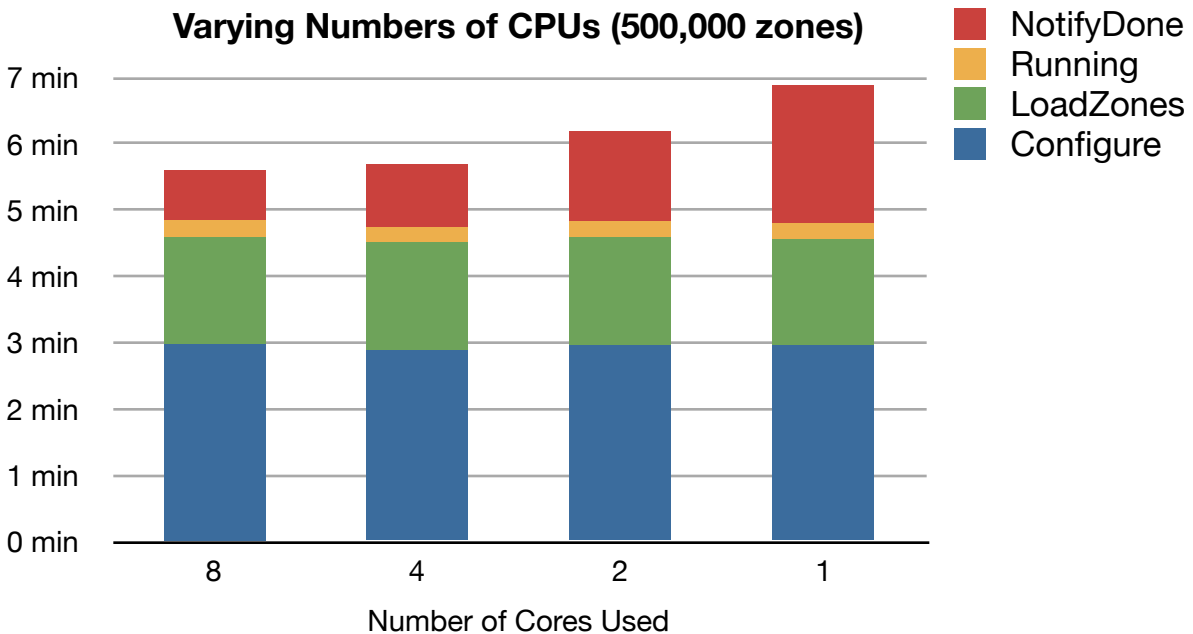
This table shows the current, fixed setting of “8” as well as some of the tuning ranges we measured. Other sizes were used as well, but omitted here for simplicity. These graphs show that a value of 1% of the total number of authoritative zones is a reasonable value for performance.

In addition to startup times, the time needed to shut down the DNS server was also significantly improved.



## Number of CPU Cores In Use

The improvement in startup times is largely unaffected by the number of CPU cores in use. In this graph, the Configure, LoadZones, and Running portions are similar, while the NotifyDone portion is improved with more CPU cores.



Current BIND 9 architecture performs a large amount of startup processing in what is called “task exclusive” mode, where only one CPU core is used. Once the server is running and answering queries, this mode is released and each CPU core can be used independently. Notices happen after this exclusive mode has ended, and will be improved because of it.

## Memory Use Increase

There is a small memory increase to allow for this improved performance, however. The memory use with 8 CPU cores is approximately 1.7% higher for the 1% scaling selection. This is further justification that the 1% is a good selection as it gives similar performance to 10%, but impacts memory use only slightly more than the 0.1% scaling factor.

## Test Machine and Configuration Details

In producing the graphs seen above, we used a standard, readily available platform with 12 GB of RAM, two quad-core Xeon processors (total of 8 cores) running at 3.0 GHz, and standard SATA drives configured without any raid or mirroring.

Each test zone was loaded from one of 1,000 different physical files. Each file was identical in its content, which included one SOA record, two NS records, and two A records.

Different sizes of zone files used and different number of physical files will affect the magnitude of the improvement. ISC welcomes any operational feedback on this improvement and will use it to further tune this and other startup related performance issues.

## Conclusion

Startup performance has been significantly improved at the cost of a modest increase in memory requirements. Results of testing with 500,000 small zones were presented. Additional measurements with 1,000,000 zones were performed by ISC but were not included here. The improvements were similar.

## About ISC

More information about Internet Systems Consortium (ISC) can be found on our web site, <http://www.isc.org> where we offer software downloads, information about our products, and information about support for our products.