

TaiRox™ Memo

Sage 300 Performance

Summary

This memo explains how small changes in a multi-user Sage 300 environment can significantly degrade or improve performance. The explanation models response time versus the number of active users, the effects of which are often misunderstood. We believe the model, while simplified for clarity, accurately describes the effect on performance as the number of users grows. **What can you do to positively affect Sage 300 performance?** See our final comments.

Model Goals


Our Sage 300 model illustrates the effects on response time for an operation that takes, on average, 1.9, 2.0 or 2.1 milliseconds in a single-user environment. That such a small difference can significantly affect performance in a multi-user setting may be surprising to those unfamiliar with queueing theory mathematics.

Model Parameters and Their Meaning

We consider from 1 to 14 SQL Server requests that occur at nearly the same time. In a Sage 300 environment, this would approximate the load for 5 times as many users, depending on a number of network factors. So, 14 simultaneous requests would approximate the load of a 70 user Sage 300 environment.

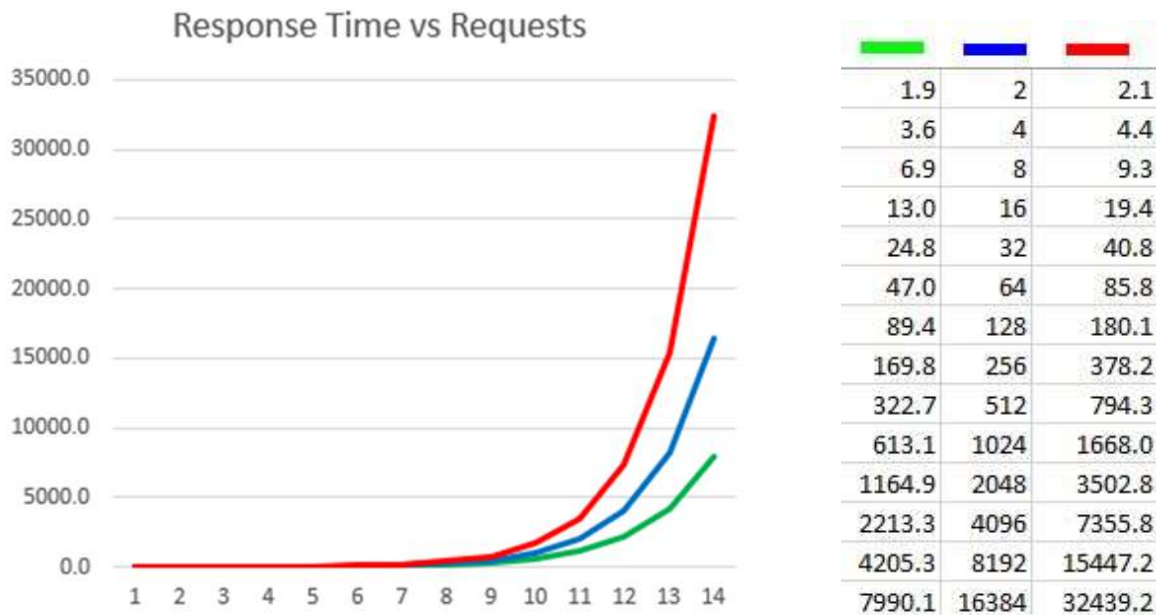
Modelling a 2 Millisecond Operation

- The chart below shows the response time in milliseconds, for 1 to 14 simultaneous requests (approximating 5 to 70 Sage 300 users).
- Up to 10 simultaneous requests (50 users), the response time is under a second - very acceptable performance.
- It may be surprising how fast response time grows to 16 seconds for 14 simultaneous requests (70 users).



Small Variations

- Increasing the parameter from 2 milliseconds to 2.1 milliseconds means the response times become unacceptable for a smaller number of users.
- Decreasing the parameter from 2 milliseconds to 1.9 milliseconds means the response times remain acceptable for a larger number of users.
- The data is shown: basic exponential progressions.



What Can You Do?

- Decrease network latency. There are well-known ways to do this. Before reaching for a quick silver bullet, measure throughput from the computer running Sage 300 to the SQL Server. There are free tools that will do a good job and leave no footprint on your network, for example Jperf and Google's xjperf. Measure throughput when the system is busy and when it is not.
- Remove anti-virus software from the SQL Server computer. SQL Server, in a Sage 300 environment, is normally NOT accessible from outside the network. So, it is NOT an entry point for an infection. TaiRox has seen dramatic performance degradation from the use of anti-virus software on a SQL Server computer. You may have to modify your compliance controls from "Malware is install on every computer" to "Malware software is installed on every computer accessible from outside the local network". You may have to negotiate with security professionals to test this change.
- Make sure that Sage 300 indexes are regularly de-fragmented. If you don't know what this means, read our memo, [Reindexing Sage 300 Tables](#).
- It may be possible, operationally, to separate multiple Sage 300 companies into 2 separate "Sage 300 networks". This will require a second SQL Server computer and probably the use of RDP servers. If this is possible, it is likely to have a huge effect.