AUTODESK®
BUILDING INFORMATION MODELING

Cost-Effective Revit® Server Deployments

Recommendations for users wishing to balance performance and hardware cost when implementing Revit Server on a project.

Introduction

Revit Server facilitates collaboration between geographically dispersed sites by replicating and caching model data at multiple points across a wide area network (WAN). In most cases it is recommended that a full “hub and spoke” configuration be deployed, where each participating site has its own dedicated Local Server connected to the Central Server. This ensures that most model data remains present within each site’s local area network (LAN) yielding optimal performance.

However, some may find deploying for optimal performance prohibitively expensive. This may be particularly true for firms featuring a large number of sites with relatively few users and relatively small hardware budgets. Our tests have revealed that in some cases administrators need not install a Local Server at each site to achieve reasonable performance. Instead, some users may be able to connect directly to the Central Server, bypassing the use of a Local Server. This can reduce the cost of deployment for Revit Server.

We explored the range of conditions where the use of this alternate configuration might be advisable, comparing the performance at various latencies to that of a reference configuration, where the user is connected to a Local Server within a LAN with the Central Server 100ms away.

Scenario 1: Excess Bandwidth Availability

In this scenario, performance for common worksharing operations was measured when using a direct connection to the Central Server across the WAN, and compared to similar operations conducted on the reference configuration. For these tests, no bandwidth limitations were imposed, short of those imposed by the networking hardware itself.

Test results indicate that connecting directly to a Central Server across a 15ms WAN yields a similar experience to that of the reference configuration. If a user attempts to connect directly to a Central Server across a WAN and latency exceeds this threshold, performance may not be reasonable and it is recommended that a Local Server be used instead.

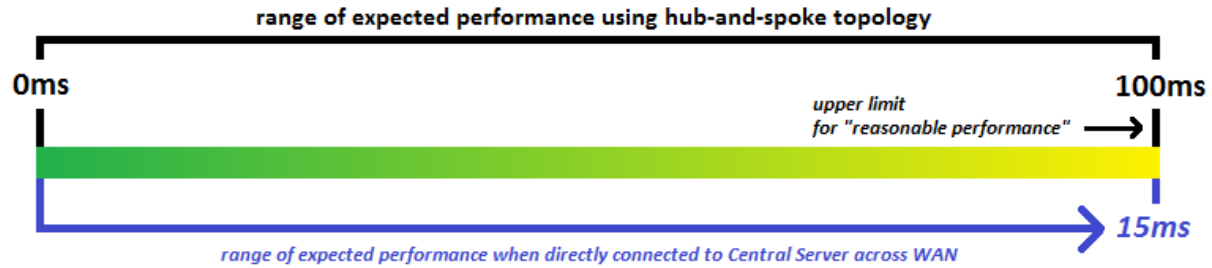

Many firms employ hardware-based data compression technology to reduce traffic across the WAN. Thus a comparable series of tests were run using Riverbed Steelhead devices to observe whether their use had any effect on the range of latencies permissible for directly connecting to the Central Server.

Our results indicate that compression dramatically expands the range of allowable latencies for directly connecting to the Central Server: using network compression, connecting directly to the Central Server across a 40ms WAN yields performance similar to that of the reference configuration.

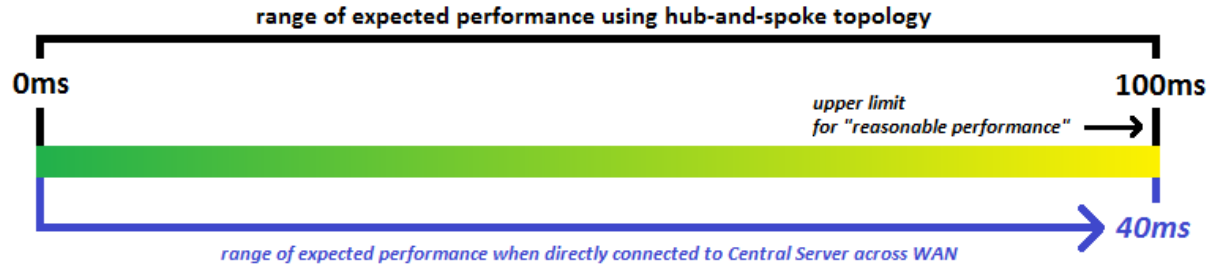

It is important to understand, however, that the degree of improvement depends entirely on the rate of compression that can be achieved in real-world use. In our testing we observed average compression rates exceeding 60%. It may be that our methodology is particularly favorable to compression. Real-world behaviors tend to be more random in nature and may not realize comparable rates of compression. Therefore we advise caution and careful testing to estimate as accurately as possible the potential improvement that can be achieved in each unique case.

As an aside, it should be noted that these tests also indicate that network compression may confer significant benefit to users accessing the Central Server via a Local Server. However, that investigation falls outside the scope of this discussion.

Scenario 2: Limited Bandwidth Availability

Internal studies indicate that Revit Server's performance degrades significantly when network bandwidth falls below 1 megabyte per second. This is especially true for worksharing operations such as Reload Latest and Sync with Central, which are generally characterized by transfers of significant volumes of data. As a result, it is generally recommended that users provide at least 1MB/sec for use with Revit Server, though some may find they need more than this, particularly if their servers are subjected to significant loading from multiple projects and/or multiple users.

For this scenario, an attempt was made to identify the range of latencies acceptable for bypassing a Local Server when operating at the minimum required bandwidth of 1MB/sec. This bandwidth rating is "real-world" bandwidth, calculated and verified by measuring the time required to move a known volume of data between two points on the network. (In live deployments, real-world bandwidth may fall significantly short of theoretical bandwidth ratings due to external conditions such as competing network traffic.) This information may be of particular interest to users operating networks where bandwidth has been intentionally constrained.

Test results reveal that when server bandwidth is constrained to 1MB/sec, the range of allowable latencies for connecting directly to the Central Server and achieving reasonable performance is quite narrow: when latency exceeds ~10ms, the time needed to complete data intensive operations such as Sync with Central exceeds the maximum allowed threshold.

The server-based workflow for collaboration moves data more efficiently than the file-based workflow, but this improvement is not enough to overcome the combined effects of constrained bandwidth and network latency when transferring data over the WAN.

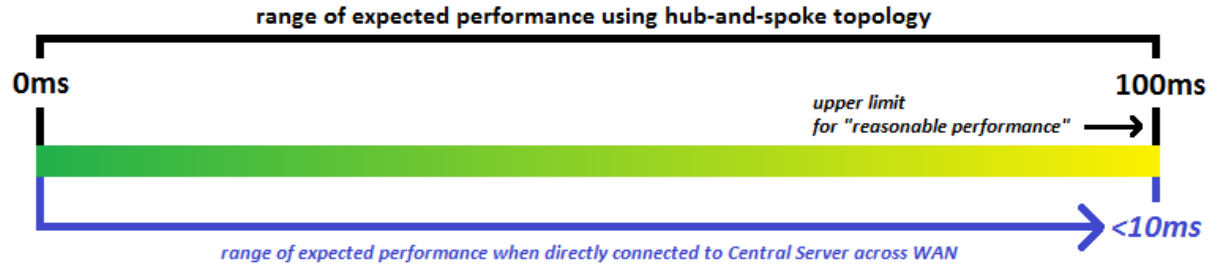

When a similar test is run employing hardware-based compression, it can be seen that this technology can mitigate the effects of limited bandwidth by reducing the amount of data that needs to be transferred. In this configuration, users can connect directly to the Central Server over a 30ms WAN and expect to have an experience comparable to that of the reference configuration.

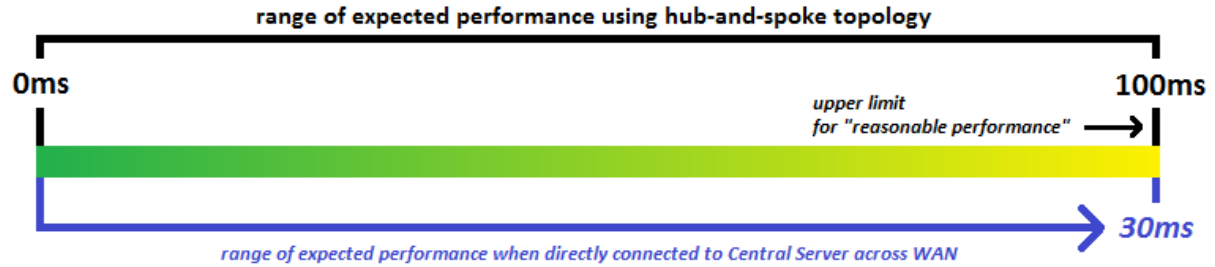

As in the first scenario, the degree of benefit realized with network compression is directly related to the rate of compression achieved when transferring project data.

Conclusion

Users collaborating over wide area networks with moderate latency and high bandwidth may find that they can bypass the use of a Local Server and connect directly to the Central Server. This potentially reduces hardware costs significantly, with only a modest penalty in application response. However, if latency is higher or bandwidth is constrained, deploying Revit Server in this fashion is inadvisable unless users are willing to accept significantly poorer application response (and therefore lower productivity) as a tradeoff for lower hardware costs.

Hardware-based network compression can be very effective at reducing network traffic. The successful use of this technology can significantly expand the range of network conditions where a direct connection to the Central Server can be used. But compression rates vary greatly from case to case, and the degree of improvement can be hard to predict. In particular, the results realized in our own tests may not be easily achieved in real-world use.