AUTODESK®
BUILDING INFORMATION MODELING

Bandwidth Requirements for Revit® Server

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

Introduction

In general, as bandwidth increases, so does Revit Server's performance. But increasing network bandwidth arbitrarily can increase costs unnecessarily. Conversely, IT managers may wish to constrain Revit Server's use of the network to accommodate other kinds of traffic.

Balancing cost, user productivity, and resource consumption is crucial to successfully deploying Revit Server. To better understand this tradeoff, tests were conducted to measure the time required to complete common worksharing operations on a 30ms WAN while adjusting available bandwidth. This document summarizes our findings.

Results and Data

Changing the amount of bandwidth available to Revit Server most strongly affects those operations that move large amounts of data across the WAN. For example, although checking out a workset accesses data on the central server, the amount of data moved is relatively small so bandwidth has little effect on performance (except in cases where it is unreasonably limited).

WAN network bandwidth	Time to checkout workset
0.193 Mbit/sec	1.51
0.386 Mbit/sec	0.96
0.772 Mbit/sec	0.73
1.544 Mbit/sec [T1]	0.65
3.088 Mbit/sec	0.65
6.312 Mbit/sec [T2]	0.63
unconstrained	0.62

Other operations with a comparably light network footprint (such as borrowing and relinquishing elements) show a similar pattern.

By comparison, operations such as synchronize with central (SWC) generally move more data across the network and therefore tend to be more strongly affected by available bandwidth:

WAN network bandwidth	Time for "typical" SWC
0.193 Mbit/sec	163.535
0.386 Mbit/sec	86.611
0.772 Mbit/sec	54.101
1.544 Mbit/sec [T1]	37.939
3.088 Mbit/sec	31.497
6.312 Mbit/sec [T2]	26.286
unconstrained	23.571

Operations that are highly data-intensive can be conceptually simplified as a basic transfer of data to or from the central server. Creating a new local copy, updating the local server cache, and syncing a very large change to the central model can all be characterized in this fashion. Here, there is a nearly linear relationship between performance and available bandwidth:

WAN network bandwidth	Time for "large update" SWC
0.193 Mbit/sec	371.55
0.386 Mbit/sec	186.12
0.772 Mbit/sec	99.09
1.544 Mbit/sec [T1]	56.22
3.088 Mbit/sec	34.41
6.312 Mbit/sec [T2]	23.35
unconstrained	15.10

Conclusion

Reasonable performance can be consistently achieved at moderate latencies if actual available network bandwidth exceeds 1Mbit per second. Further increases in bandwidth will increase performance, but for some operations there may not be a commensurate increase. Note that actual bandwidth often falls short of theoretical bandwidth.

Those who need to accommodate frequent simultaneous server access may wish to make still more bandwidth available, as these environments are likely to witness frequent parallel data transfers, which must share network bandwidth. Note that there is a distinction between “simultaneous access” and “simultaneous connections.” Revit users who are connected to Revit Server are not necessarily moving data to or from it, so they are not actually consuming network bandwidth. Thus, when predicting bandwidth consumption, it is as important to consider user behavior is it is to consider user count.

Those operating at high latencies may also wish to make more bandwidth available, as low bandwidth and high latency can have a negative compound effect on the usability of Revit Server.

Those who cannot satisfy our minimum recommended bandwidth requirement may wish to employ network traffic compression technology such as Riverbed’s Steelhead devices, as compression can be particularly helpful under low bandwidth conditions.