

Panzura White Paper

---

# CAD Theory of Operation & Best Practices

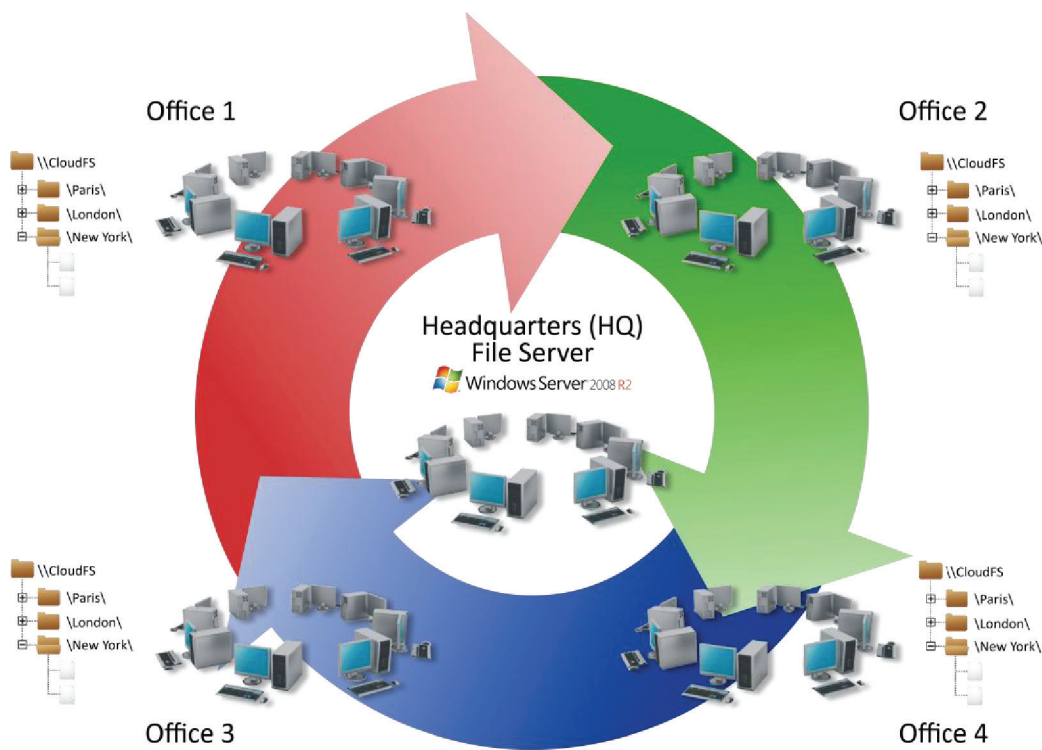
Panzura's game-changing Distributed Cloud File System technology finally brings the full power and benefits of cloud storage to enterprise customers, helping to break the unending on-site storage expansion cycle while eliminating islands of storage that inhibit cross-site user interaction and productivity and real-time data protection. Panzura's Distributed Cloud File System enables CAD design files to be shared over the WAN as if they were on the LAN. This white paper explains how.

## Executive Summary

Over the last 30 years, CAD modeling has grown from a single user CAD design product running on a microcomputer to a complex CAD design environment with hundreds of millions of seats sold with multiple derivative products spanning multiple market verticals. Architecture-Engineering-Construction (AEC) firms large and small struggle in providing adequate infrastructure to allow its users sufficient performance to satisfy even the most basic functionality, such as file open and save. Panzura Freedom Flash Cache appliances offer high performance and a cost-effective solution that seamlessly integrates with CAD files. By leveraging Panzura, AEC organizations – with users dispersed across multiple geographies – can share CAD files as though the data is right next to them. This document describes a set of best practices, along with a contextual background, on how everything works.

## Background

CAD design software is used by leading architecture, engineering, and construction firms spanning several different markets. This white paper describes common workflows used by CAD engineers dispersed across multiple geographies within the AEC industry. It applies to all desktop versions of CAD modeling files and derivative products, including AutoCAD Map 3D, AutoCAD Civil 3D (built on Map 3D), SOLIDWORKS, and a variety of other popular products, including Revit and Bentley MicroStation, where similar types of file operations are performed. It then goes into details of the common problems experienced by these users and how Panzura Freedom Flash Cache appliances resolve these problems both at a high-level engineering view and the detailed IT view.



In this paper we will use AutoCAD for illustrative purposes. However the same principles could apply to any CAD application.

## File Open Problem

One of the most common problems CAD users voice is the time it takes to open and save a file remotely. It frequently takes 15-30 minutes to open an AutoCAD DWG file when it only takes 5-10 seconds for the same operation locally. This problem, in fact, is so common that the most popular CAD training sessions go to great lengths in providing a strategy to what is commonly known as the “File Open Problem”. Dramatic changes in workflows are suggested in order to improve performance, yet these changes rarely fix anything.

The first reaction IT professionals managing AEC firms have when facing this problem is that there must be a WAN bandwidth problem. Many firms were limited to lower than T1 (1.54Mbps) speeds during the early and mid-2000s, and the explanation sounded logical then. AEC firms typically function under very tight budgets and WAN increases were a luxury most couldn't afford just five years ago. So the assumption was widely accepted without thorough validation, as the verification process was impractical. But in recent years, WAN bandwidth prices have collapsed and many firms have been able to upgrade to DSL/cable/broadband/3-4G speeds of 6Mbps or higher at a fraction of T1 pricing. Many can verify this by simply enabling 4G tethering on their smartphones or visiting a Starbucks with much higher than T1 bandwidth. Yet the problem persists in every remote location and gets worse based on distance.

A close analysis shows that opening typical CAD files, ranging in size from 1MB to 100MB, still take 15-30 minutes to open and save from remote locations. Yet in theory, those very files should take no more than two minutes if bandwidth was the issue. So clearly the problem warrants further analysis.

## Main Culprit

In short, the culprit is latency. The further one is from the source of the data, the higher the latency and the longer AutoCAD files take to open and save. It is a basic speed of light problem. As AutoCAD DWG files get more and more complex, they embed more and more objects and each object takes longer and longer to open and save based on latency. The size of each object is an issue only when bandwidth is under 1Mbps. That may have been common 5-10 years ago but is no longer the case. The main delay is the sheer number of objects (complexity of DWG file).

## Referenced Drawings (Xrefs)

AutoCAD has many different types of objects that increase load and save time as latency increases. One of such types used universally are objects called Referenced Drawings or Xrefs. AutoCAD user guide describes the Xrefs as:

You can attach an entire drawing file to the current drawing as a referenced drawing (xref). With xrefs, changes made in the referenced drawing are reflected in the current drawing. Attached xrefs are linked to, but not actually inserted in, another drawing. Therefore, with xrefs you can build drawings without significantly increasing the drawing file size.

By using referenced drawings, you can

- Coordinate your work with the work of others by referencing other drawings in your drawing to keep up with the changes being made by other designers. You can also assemble a master drawing from component drawings that may undergo changes as a project develops.
- Ensure that the most recent version of the referenced drawing is displayed. When you open your drawing, each referenced drawing is automatically reloaded, so it reflects the latest state of the referenced drawing file.
- Keep the names of layers, dimensioning styles, text styles, and other named elements in your drawing separate from those in referenced drawings.

Merge (bind) attached referenced drawings permanently with your current drawing when the project is complete and ready to be archived

Xrefs are commonly used across all AutoDesk product lines and one can logically see why they are especially popular in projects that are edited by many CAD designers across multiple sites. Xrefs reduce redundant work, ensure the latest and greatest is always used and allows inheriting features within an object while providing customization of the object specific to the project in question.

## Technical Detailed View of Loading Xrefs

Now that we know that Xrefs are one of many objects that are slow to load remotely, let's take a technically detailed view of what exactly happens when we load a DWG file with Xrefs. Customers interested in free analysis can email [cad-analysis@panzura.com](mailto:cad-analysis@panzura.com). Panzura helped analyze DWG file opening provided by several AEC firms under NDA and they were all virtually identical in the way data moved from source to remote. Below is an analysis of a simple 1.5MB Civil3D file courtesy of C&S Companies. The file was opened over a typical Windows File Share mount from a Windows 7 client to a Windows 2008 server. The protocol being run across the client and server was SMB2:

### Opening of Civil3D file:

File size: 1.5MB  
 Purpose: Building design  
 Total read data packets: 1,004  
 Total file opens and locks: 4,508  
 Total file closes and unlocks: 4,206  
 Total directory browsing: 5,385  
 Total network packets: 15,103 (sum of above)

C&S headquarters is based in Syracuse, New York. We opened this file from two different remote locations, both with models 6Mbps WAN bandwidth (DSL). The two locations were San Diego and Phoenix. Results are as follows:

### San Diego, CA<----->Syracuse, NY

Round trip latency: 86msec  
 WAN bandwidth: 6Mbps  
 Time to open directly: 22min  
 Time to open over Riverbed cold: 24min  
 Time to open over Riverbed hot: 20min  
 Time to open over Panzura: 8 secs

### Phoenix, AZ<----->Syracuse, NY


Round trip latency: 66msec  
 WAN bandwidth: 6Mbps  
 Time to open directly: 16min  
 Time to open over Riverbed cold: 17min  
 Time to open over Riverbed hot: 15min  
 Time to open over Panzura: 8 secs

As one can see, the bulk of the delay is spread across three types of packets; "Open and Locks", "Close and Unlocks" and "Directory Browsing". The data read made up less than 10% of the total delay in opening the file. "Open and Locks" is when an application tries to open and lock the file for read or write. This is common across all desktop applications such as Microsoft Office, AutoCAD and many others. "Close and Unlocks" are the exact opposite of "Open and Unlocks". "Directory Browsing" are issued when an application tries to browse a directory looking for a specific file. All three types are very common in desktop applications. The reason there are so many is because AutoCAD DWG project files that are worked on by multiple users across multiple locations must include a certain amount of shared objects. And the opening those files results in a lot of network packets.

One may ask why there isn't technology that aggregates or optimizes away these packets. The answer is simple. They cannot unless the technology includes the filesystem. These are basic filesystem commands that must be issued and answered by the filesystem. If a cache or WAN optimization technology were to optimize a lock away, then the file or data being locked may be modified by multiple entities across geographies resulting in conflicts. This is indeed what many AEC firms have experienced when using technologies such as Microsoft DFSR. This is why no technology can optimize this away. The only solution to this problem is by having a distributed filesystem so the replies are always atomic and issued directly at the remote site.

## AutoCAD with Panzura Cloud Storage Controllers

Panzura recommends a hub and spoke configuration for AEC customers. Please review [panzura.com/resources/white-papers/](http://panzura.com/resources/white-papers/) "Panzura Distributed Cloud File System – Overview" for an in-depth analysis of how the Panzura filesystem functions. The diagram below describes a typical hub and spoke topology:



## Benefits of Using Panzura Freedom Flash Cache

- **Instant locking of files** – this provides instant file open, lock, unlock, or close from both headquarters and remote locations. Please review the white paper on distributed locking at [panzura.com/resources/white-papers/](http://panzura.com/resources/white-papers/) “Distributed File Locking”. This is the main benefit for CAD file open and save speedups.
- **Seamless Windows AD/CIFS** – works exactly like your Windows 2008 R2 file server from an end-user standpoint. It’s completely seamless and users mount file shares just as they do and administrators manage the file server and joining AD for security just as they’ve been doing for the last decade.
- **Global Filesystem** – this provides one single Windows SMB mount with uniform global view everywhere. Users no longer need to mount remotely or find where each file/folder is located.
- **Cloud Reliability And Backup** – data is replicated and distributed for performance and reliability across global cloud storage providers that are used for moving and storing some of the largest and most sensitive workloads. Please visit [www.panzura.com](http://www.panzura.com) for case studies of how Fortune 500 companies use the product. Panzura Freedom appliances scale for large and small enterprises.
- **Cloud Distribution** – With Panzura, the authoritative copy of the data is securely stored within the cloud, so that all data read by branch offices is read directly from the cloud and is not requested from the HQ based storage. This means that if 10 branch offices request the same HQ-based 1GB file, HQ does not need to transmit 10GB worth of the same 1GB data out of its WAN link. Instead, the 1GB file is transmitted once to the cloud and all branch offices retrieve it from the cloud service provider, giving a drastic optimization of the HQ bandwidth.
- **Global snapshots** – data can be configured to be globally retained hourly, weekly, monthly, and yearly – up to 10,000 snapshots – so backups and archives into the cloud can be done from one central location on the same platform.

## Summary

Panzura Freedom Flash Cache, when used as a Windows SMB file share for CAD projects, provides instant file access, security, and seamless integration across multiple geographies. DWG files that used to take 20-30 minutes to open and 5-10 minutes to save now can be opened and saved in under 5 seconds. Results have been verified across both small and large AEC enterprises ranging from 5-10 users in 3-4 sites to thousands of CAD engineers across hundreds of sites. WAN optimization and other types of WAN caching products are great when WAN bandwidth (<1Mbps) is a problem. However, at today’s WAN bandwidth prices and availability, it rarely is. The problem is a filesystem latency issue. Customers are encouraged to evaluate Panzura Freedom Flash Cache appliances to experience the benefits of a distributed filesystem and its benefits when it comes to CAD files.

## About Panzura

Adding over two petabytes of enterprise storage under management per month, Panzura is a market leader in enterprise hybrid cloud storage. Panzura's hybrid cloud storage software seamlessly combines the economics, capacity, and business model of cloud storage with the flexibility, performance, and features of enterprise storage. Using Panzura, enterprises shift from an antiquated storage status quo model into the cloud and unlock the power of data. Organizations like the Department of Justice, Milwaukee Tool, American College of Radiology and Chevron use Panzura for active archive, hybrid cloud NAS, and cross-site collaboration. Please visit [www.panzura.com](http://www.panzura.com) for further information.



Panzura, Inc. | 695 Campbell Technology Pkwy #225, Campbell, CA, USA | 855-PANZURA | [www.panzura.com](http://www.panzura.com)  
Copyright © 2017 Panzura, Inc. All rights reserved. Panzura is a registered trademark or trademark of Panzura, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.