

## TECHNICAL PRODUCT SPECIFICATION

FUNCTIONALITY	ZEBCLIENT SUPPORT
<b>Deployment Modes</b>	<ul style="list-style-type: none"> <li>Standalone, Cluster and Distributed Cluster</li> </ul>
<b>Supported Infrastructure for Deployment</b>	<ul style="list-style-type: none"> <li><b>OS:</b> all major Linux OS distributions, (Ubuntu, Redhat, Debian), suggested OS: Ubuntu 20.04 LTS or later</li> <li><b>Memory / Disks</b> (hot and warm tiers): 3 units in total. PMem, NVMe or SSD.</li> <li><b>Processes run on:</b> VMs, OS on raw metal, light virtualized containers or Kubernetes</li> <li><b>Cold storage:</b> S3 cloud account</li> </ul>
<b>Input and Output Data Interfaces</b>	<p><b>S3:</b></p> <ul style="list-style-type: none"> <li>HeadObject, CopyObject, PutObject, GetObject, DeleteObject</li> <li>PutBucketVersion, GetBucketVersion, GetBucketLocation</li> <li>ListVersionedObjects, ListObjects,</li> <li>DeleteMultipleObjects</li> <li>HeadBucket</li> <li>PutBucket, DeleteBucket, ListBuckets</li> </ul> <p><b>Filesystem:</b></p> <ul style="list-style-type: none"> <li>Reading and writing files</li> <li>Seeking while reading / writing</li> <li>Append mode while writing</li> <li>Truncating at arbitrary offset</li> <li>Removing files</li> <li>Creating and removes directories</li> <li>Listing directories</li> </ul>
<b>Internal Format and Communication</b>	<p><b>Internal Format:</b> Data stored as erasure coded objects (shards)</p> <p><b>Internal Communication:</b> gRPC between hot and warm tiers</p>
<b>Tiering</b>	Storage aware class-of-service (COS). Note that tiering is a logical construct and the number of tiers can be arbitrary. However, each tier should represent a homogenous class from a performance perspective. A higher tier must be faster and smaller than a lower tier,
<b>External Cold Storage</b>	External to ZebClient via S3 The ZebClient configuration file will list the desired endpoint and protocol, e.g., connection to public S3 compatible cloud.
<b>Data Read/Write Boost</b>	The tiering function will place hot data on the level with fastest access and lowest latency possible, typically PMem when present. R/W from ZebClient (PMem, NVMe SSDs, SATA SSDs) is boosted with data stored on multiple parallel disks ("striping" effect).
<b>Redundancy</b>	<ol style="list-style-type: none"> <li>Erasure coding using redundancy shards</li> <li>Support for shard distribution to separate physical hardware units</li> </ol>
<b>Configuration Interface</b>	CLI and configuration file
<b>Metrics &amp; Monitoring</b>	Embedded Prometheus function to collect general metrics. Displaying general system data, e.g., CPU, memory, disk, network as well as application specific data.



### The ZebClient Software

ZebClient optimizes the utilization of modern memory technology to provide sustainable, redundant hyper-speed access to your data. It is a data access point software that operates by localizing cloud data to memory and by disaggregating compute from storage. This allows for a solution that speeds up, simplifies, secures and protects the information under its control and makes local storage redundant.

ZebClient offers a new versatile, scalable, hyper-performant and global data accessibility solution designed based on cutting edge technology. From its unique object data design and the capability to handle both cloud and legacy data. ZebClient also bridges the gap between old and new IT technology.



### Support

Online support via:  
<https://support.zebware.com/hc/en-us>