

### CAStor®

The Market Leading  
Object Storage Engine

Massive scalability and  
future-proof accessibility  
of unstructured data

The foundation of the Caringo® Object Storage Software Platform is CAStor®, software that enables massive scalability and future-proof accessibility of unstructured data. Utilizing any combination of standard x86 server hardware, any size hard drive and any type of storage, CAStor pools server and networking resources combining them with advanced management, optimization, and self-healing technology. The result is a unified storage system utilizing a single namespace that can scale to billions of objects and petabytes per location with superb performance, and guaranteed data protection.

#### Performance – responsive even at petabyte scale

**No file system limitations** — Universally unique IDs stored in a flat address space eliminate the complexity of hierarchical file system architectures. There are no limits on quantity of files and no capacity limits. Achieve the performance of primary storage for content of all types, large or small.

**Instant object identification** — Universally Unique Identifiers (UUIDs) are assigned to each object for life. Because the UUIDs are location independent, applications do not need to know where a specific piece of content is, yet they can find it anywhere in a node, cluster or remote cluster – instantly.

**Simple, adaptive, symmetric software** — CAStor uses symmetric clusters to evenly distribute processing across all nodes. Performance is continuously optimized using adaptive algorithms to identify the right nodes, at any point in time, for requested operations.

**Plug-and-play throughput** — Insert a new node into the cluster and it is automatically recognized and its capacity is pooled with the overall cluster capacity with no downtime and no impact to applications.

#### Efficiency – optimize management and operational resources

**Easy to install** — CAStor software can boot from a USB key or can PXE boot from the network.

**Self-managing** — Symmetric cluster architecture evenly distributes processing across all nodes. Operation requests go to the nodes best able to handle them to optimize performance. CAStor not only distributes CPU loads but also organically balances storage itself. There are no specialized nodes which means there are no bottlenecks. And if there is a better node on which to store content, CAStor will move it in the background to ensure automatic evenly distributed content migration.

**Seamlessly upgrade and retire hardware** — Add new server nodes to the cluster then simply retire the old server and remove from the cluster. Upgrade hardware technology at your own pace and leverage price performance curve of new technology as soon as it becomes available.

**Integrated power conservation** — Darkive™ is a patent-pending adaptive power conservation technology that monitors storage system operations and automatically spins down disks and reduces CPU utilization.

### Recommended Hardware Requirements

CAStor runs on a cluster with a minimum of three (3) nodes.

**CPU:** 64-bit, Intel Xeon, AMD Athlon64, or equivalent

**Hard Drive:** Standard SATA with 1 or more drives per node

**Memory:** 1GB RAM per disk drive

**Network Interface:** Dual Gigabit Ethernet

### About Caringo

To learn more about Caringo and our products visit our website or e-mail [info@caringo.com](mailto:info@caringo.com).

#### Caringo, Inc.

6801 North Capital of Texas Highway  
Building 2, Suite 200  
Austin, Texas 78731  
512-782-4490

[www.caringo.com](http://www.caringo.com)

### Protection – built in security, integrity and compliance

**Continuous integrity checks** — CAStor's Health Processor continuously monitors data integrity and cardinality (number of replicas) and will automatically heal any degradation or non-conformity. Metadata life point rules are also checked and enforced.

**Content integrity seal** — CAStor establishes content authenticity for evidentiary purposes by creating and storing content integrity seals. Seals are requested and stored by the application at time of write on a per object basis, computed on the content only, not the metadata.

**Patented, transparently upgradeable hash** — Update your hash protection without impacting your applications and stay up-to-date with most current hashing algorithms before they become compromised.

**Object-level retention** — Retain content for a specified period of time to meet regulatory or governance policies. Set retention time on a per object basis and the system will ensure it remains available and unchanged until it expires.

**WORM storage** — CAStor addresses regulatory mandates that content be stored on non-erasable, non-rewriteable media. True to Write-Once-Read-Many (WORM) storage media, once an object is written it cannot be changed or deleted

### Scale – no limit in capacity or object count, ubiquitous access, and breakthrough economics

**Quickly meet any demand** — Plug-and-play scale by simply inserting one, or more nodes into the cluster. Additional capacity is automatically available with no manual provisioning required. A single IT administrator can easily manage multiple petabytes in a single cluster.

**Unlimited flat address space** — A single, flat 128-bit address space enables extremely efficient replication and migration. It provides massive and seamless scalability, long-term archiving and the flexibility to join or split storage infrastructures, all without changes to the application.

**HTTP interface** — Applications and devices interface to CAStor using industry standard HTTP 1.1. This enables quick development and the ability to utilize existing HTTP client libraries in use throughout the industry.

**Breakthrough economics** — Storage clusters can be built on any X86 server. Use any combination of vendors and generations. Plug in server in and CAStor will instantly provision it.

### Control – robust management simplified by automation and standards based tools

**Flexible object naming** — CAStor offers unrivaled flexibility in naming supporting user defined named objects and system assigned unique IDs.

**Granular control per object** — Metadata is the DNA of digital content. It describes where an object came from, what's in it, who can access it, how it's protected, replicated, and distributed, and how long it is destined to live.

**Self-optimizing** — Even after an object is initially stored CAStor monitors storage utilization efficiency and the resulting performance automatically moving or replicate data in order to maintain balance and responsiveness in the system.

**Self-healing** — Policies set for each object are checked and enforced by the health processor. Should something fail, the system automatically recovers.

**Web-based storage administrative console** — Manage, monitor and control the entire cluster from a web-based management and administration user interface.

**Syslog server support** — Alerts can be sent to a Syslog server for standard reporting, diagnosing alerts or establishing audit trails.

**SNMP and NTP support** — Leverage an existing systems management platform to manage/administer the cluster as another element in the network as well as support external network clock synchronization.

**Intelligent Platform Management Interface (IPMI) support** — CAStor leverages IPMI to monitor the health of the CPU, disk drives and other components with the cluster and alert administrators if a potential problem emerges.