Distributed Storage

We've always had to distribute and replicate data, whether to protect ourselves from data center outages or to run offices in multiple regions. This is even more relevant in the cloud as computer systems are becoming increasingly distributed. One of the most difficult challenges of distributed storage is coordinating how data is accessed and updated.

Traditional Approaches

To protect against hardware failures and outages, we often replicate data on multiple devices and locations. Should the "master" location fail, we can immediately fail over to the other location. Of course, if we lose contact with the secondary location at one point, we'll need to resynchronize the secondary storage when contact is re-established to bring it up to date before we can consider it fully valid.

For more complex needs, we can also use solutions such as Ceph, HDFS, and erasure coding to spread data to many locations and devices. These solutions require careful planning and a setup of controllers or clusters to manage, making them rather complex to deploy.

Finally, ledgers such as blockchain can also decentralize storage, but they require powerful computations and intensive peer-to-peer coordination to validate. Each node also needs to copy large amounts of data onto local databases or storage to be helpful, making these particularly slow and expensive solutions.

While we have certain technologies to distribute data on multiple locations, they are either limited, complex to set up, expensive, or significantly slow.

Current Limitations

While conducting COSNIM scientific research trying to find better ways to distribute data securely in the cloud, it became increasingly apparent that the primary reason why decentralized storage management is so complex and challenging is that we're still basically managing data as if they were pieces of paper. For example, backups are essentially the electronic equivalent of taking photocopies of pages and then putting them in archival boxes; replication is the equivalent of sending photocopies to a remote office, tracking and resending copies when the original is altered. This is why distributed systems require servers, clusters, and careful peer-to-peer or peer-to-server coordination to operate safely and maintain integrity.

COSNIM Technology

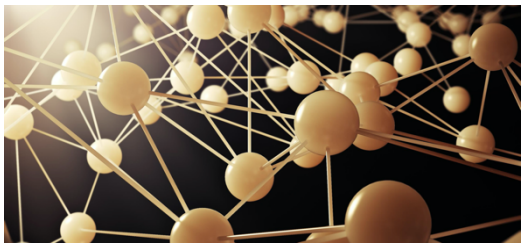

Instead of trying to leverage or enhance existing technologies, COSNIM takes a dramatically different approach. All fundamental pieces of data in COSNIM are treated as self-governing elements, forming free bonds between themselves to build a self-supporting structure, much like atoms bind themselves together to form large, complex molecules. COSNIM relies entirely on this unique structure to operate, without needing servers or a central authority, just as molecules don't need any external physical support to exist or interact with their environment.

When used for file storage, COSNIM breaks up the file's contents into fragments and binds them together with metadata and control information. Elements are packaged into one or more Capsules (which may contain other unrelated elements) and bonded in turn with other capsules in the Continuum. This creates a flexible mesh of interrelated elements and capsules, similar to how large molecules are built from smaller molecules and individual atoms. COSNIM uses highly efficient processes and proprietary algorithms to manage these bonds, allowing capsules and their data to be freely stored and distributed anywhere in the cloud or on local devices, on any number and type of storage, without tracking their physical location.

The Continuum

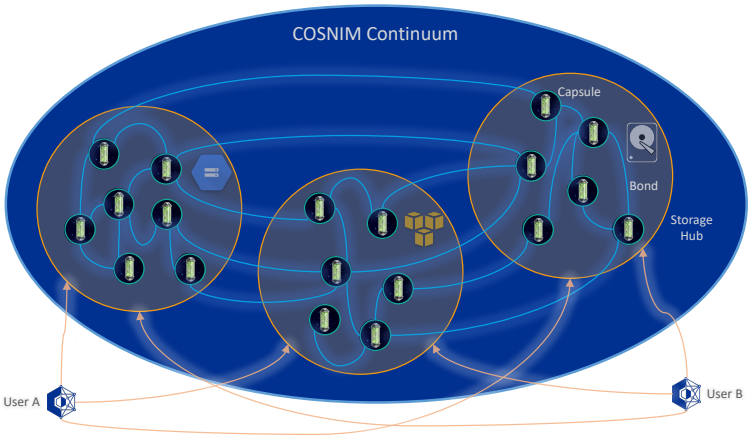

This unique, self-supporting structure is called a COSNIM Continuum. Users connect to a Continuum directly through the storage hubs where capsules are physically stored. There are no central servers, clusters, or peer-to-peer communication. Everything is held together and managed entirely through capsules and the bonds that form the Continuum. When data is updated, instead of physically replacing or updating storage units such as blocks or files, COSNIM creates one or more new capsules that slightly alter the Continuum's mesh to integrate the latest data, leaving all other components intact. There is no central authority or fixed location for any of the data; any storage hub or capsule can be used by COSNIM to explore and access the data. This highly adaptable structure and decentralized design are what give COSNIM many of its unique capabilities.

Time-Travel

COSNIM Time-Travel is an extremely efficient replacement for backups, snapshots, versioning, and journaling. Since updates are recorded as slight alterations to the bonds that hold the Continuum together, COSNIM can instantly travel through the mesh to

access live data and also any previous state. This gives Time-Travel instant and continuous access to all changes that occurred in a Continuum, without ever taking copies of data or actively tracking individual changes as traditional technologies need to do. Since Time-Travel is driven directly by how the Continuum is natively structured, it does not incur any additional processing overhead, either when data is produced or when previous states are examined.

Data Migration

Because bonds are location-independent, Capsules can be freely moved from any storage hub to any other without coordination or notification. This provides an extremely easy and flexible data migration tool, allowing data to be transparently moved from one location to another, in cloud or enterprise storage, live, even while updates are in progress.

COSNIM's unique technology enable powerful tools such as Time-Travel continuous data protection, transparent migration and highly resilient asymmetric replication.

Asymmetric Replication

To increase availability, capsules may also be freely replicated from any storage to any storage in any fashion. Replication does not need to follow any mirroring or pre-determined replication pattern as with traditional technologies, although users can do this if they want. Capsules can be replicated in any variable number of copies anywhere there's sufficient storage. During outages, COSNIM automatically and instantly accesses capsules available anywhere in the Continuum without triggering failovers. When a storage hub becomes available again, its capsules are immediately re-integrated into the Continuum, even if the Continuum was updated, eliminating all traditional resynchronization vulnerability windows. COSNIM calls this "Asymmetric Replication," it is a potent tool that ensures maximum availability and flexibility.

Distributing data storage in multiple locations, especially in the cloud, is often difficult, complex, and costly. COSNIM's proprietary technology allows users to freely distribute data anywhere in the cloud and enterprise storage with extraordinary ease, along with highly advanced capabilities such as Time-Travel.

Https://www.cosnim.com
info@cosnim.com
+1 888 212-3527