

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. In the cloud, this is even more relevant as computer systems are becoming increasingly global. One of the most difficult challenges of distributed storage systems is coordinating how data is accessed and updated.

Traditional Approaches

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


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 do require careful planning and a setup of controllers or clusters to manage, making them much more complex to deploy.

Finally, ledgers such as blockchain are also able to 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 useful, making these particularly slow and expensive solutions.

While we do have certain technologies at hand 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 difficult 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 if the original is altered; servers are the equivalent of comptrollers, which carefully coordinate and centralize the view of the data; even journaled filesystems and blockchain are based on basic ledger concepts, albeit with a few twists.



COSNIM Technology

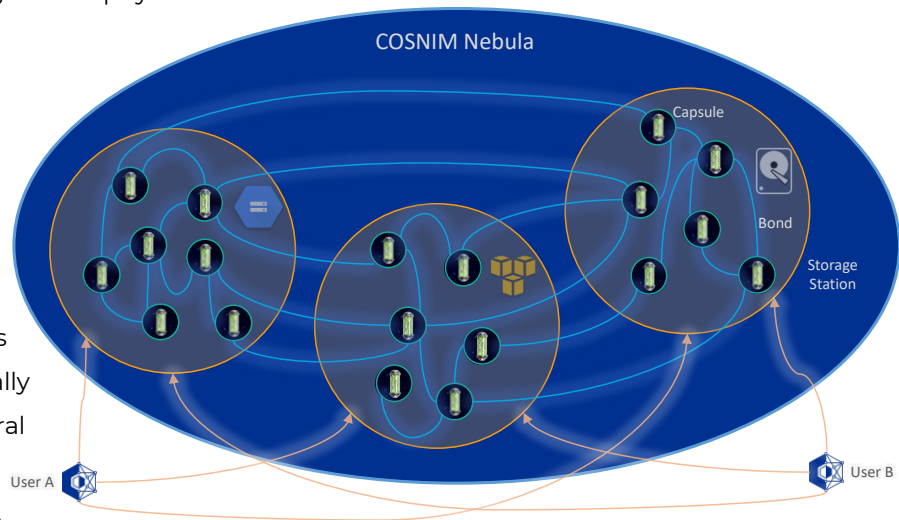


Instead of trying to leverage or enhance existing technologies, COSNIM took a dramatically different approach. All fundamental pieces of data in COSNIM are treated as self-governing elements, forming free bonds between each other to build a self-supporting structure, much the same way atoms bind themselves together to form large, complex molecules. In the same way, COSNIM relies entirely on data elements and capsules for structure, without the use of servers or central authority, just as molecules don't need any external physical support to exist or interact with their environment.

When this is applied to file storage, COSNIM breaks up the file's contents into fragments and binds them all together with metadata and control information, which are also elements. Elements are then packaged into one or more capsules (which may contain other unrelated elements) and bonded in turn with other capsules in the Nebula. This creates a mesh of inter-related elements and capsules, similar to how large molecules are built from smaller molecules and individual atoms. COSNIM uses highly efficient patented 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 any tracking of their physical location.

The Nebula

This unique, self-supporting structure is called a COSNIM Nebula. Users connect to a Nebula directly through the storage stations 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 Nebula. When data is updated, instead of physically replacing or updating actual storage units such as blocks or files, COSNIM creates small new elements and ever so slightly alters the mesh to integrate this new data, leaving all other components intact. There is no fixed location for any of the data; any storage station and capsule are fully capable of carrying everything. 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 small alterations to the bonds that hold the Nebula together, COSNIM can instantly travel through the mesh not just for live data, but also to examine all previous states. This gives Time-Travel instant and continuous access to all changes that occurred, without ever taking copies of data or actively tracking changes as traditional technologies need to do. Since Time-Travel is driven directly from the way the Nebula is natively structured, it is essentially free and does not incur any additional overhead, either when data is produced, or previous states are examined.

Data Migration

Because bonds are location independent, Capsules can be freely moved from any storage station to another without any coordination or notification whatsoever. 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. 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 anywhere there's sufficient storage. COSNIM will transparently use capsules available anywhere else in the Nebula during outages. When the storage becomes available again, its capsules are immediately re-integrated, eliminating all traditional resynchronization vulnerability windows. COSNIM calls this "Asymmetric Replication", and it is a highly powerful tool to ensure maximum availability.

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.