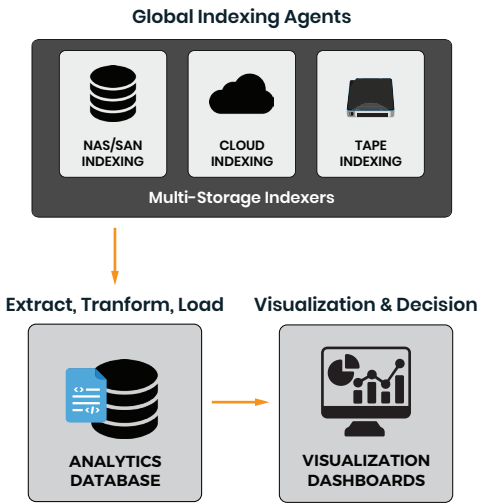


DATA VISIBILITY SERVICES

Powers visibility across multiple sites and multiple storage pools (including hard drives, cards, NAS, SAN, Object, Cloud and tape) enabling key decision makers to drive better "Data Decisions".



DNAfabric implements a "Big Data" pipeline to process vast quantities of metadata to drive customizable dashboards. The metadata is processed in phases, where end user data across different storage pools (including hard drives, cards, NAS, SAN, Object, Cloud and tape) is indexed and monitored. This metadata is then transformed using JSON schemas customizable to an end user's environment. Once transformed, it is loaded into a global analytics database. Live connectors allow customizable dashboards to highlight numerous aspects of data usage, growth, SLAs.

CHALLENGES

Growing Assets, Numerous Storage Options, Cloud Initiatives

The quantity of data continues to grow. Additionally, this growth is not across a homogenous and managed storage environment. Instead, users are grappling with a number of storage options from field hard drives, LTO tapes, Nearline storage, primary storage and cloud. This results in an unmanaged storage environment.



DNAfabric: TRANSFORMATIVE DATA MANAGEMENT AND DATA DECISION PLATFORM


One Platform, Multiple Data Services

In the world of astronomical data growth, distributed data pipelines, and growing storage options, DNAfabric is a new way to look at large scale data management. DNAfabric is an easy to deploy platform designed to address gaps in any data management strategy. It can be deployed without "rip and replace" efforts and provides a simple subscription model, offering numerous data management services. These services power flexible data movement, tighter data governance and improved data-decision making ability.



DATA MOBILITY SERVICES

Flexible services replace numerous point products for data backup, archive, synchronization, sharing and collaboration. Designed to work across both on-premise, cloud and hybrid workflows.



Multi-Direction Mobility

- NAS/SAN ↔ NAS/SAN (Local, UDP)
- NAS/SAN ↔ Cloud/Object
- NAS/SAN ↔ Tape

On-Prem or Cloud

Deploy services on physical and virtual machines across on-premise, cloud and remote sites.

Local, UDP, Object Acceleration

- Local Acceleration (Multi-threading, SSD Caching)
- UDP Remote Acceleration
- Cloud/Object Acceleration (Multi-part upload, Multi-threading)

DNAfabric mobility services is powered by synchronization agents that can mobilize data across storage pools on-premise (NAS, SAN, online, Nearline, LTO), remote sites and cloud storage. This allows users to replace numerous point tools with a single comprehensive platform.

Ad-Hoc Data Management Tools

To manage the data growth across a diverse set of storage options, users are deploying data management tools on a more ad-hoc basis. For example, field workflows are deploying offload utilities for camera card offloads. In Post, pipelines are using data management tools across primary, nearline and archive storage. And yet another set of tools are being deployed for remote transfers and cloud uploads. This results in environments that have many unmanaged tools.

Unmanaged and Un-Secured Processes

This is resulting in data environments that are largely unmanaged and poorly secured. Numerous databases, lack of accountability, poor reporting and statistics are some of the characteristic traits of any environment today.

CIOs, CFOs, Heads of IT are making decisions on poor insights resulting in increased costs and environments more prone to data loss

DNAfabric: Powerful tools for next generation of data pipelines



Powerful Data Analytics Dashboards Drive Smarter Decisions for IT Managers and C-Level Executives

Current insight tools provide "Finder/Explorer" style directory and utilization statistics. While these tools are "good enough" for gathering basic information, they provide minimal insights into how data is used, protected and scaled. On the other hand, DNAfabric's Data Visibility services is a data visualization toolset designed for C-level executives and heads of departments enabling real time tracking of data for improved provisioning, utilization, spending and protection.

High Performance Data Movers Accelerate Local, Remote and Cloud Transfers

DNAfabric is powered by data synchronization technology optimized for a number of deployment scenarios. DNAfabric accelerates local disk to disk transfers between any shared storage. In local transfers, it uses multi-threading, variable buffers and SSD caching to optimize transfers. Over remote links, DNAfabric employs UDP acceleration with multi-threading to achieve full link utilization. For transfers to cloud object stores (e.g. AWS S3, Azure Blob, WD Activescale and more), DNAfabric employs multi-part upload and multi-threading.

Drives IT Style and Media Intelligent Pipelines

DNAfabric is perfectly suited for IT style workflows designed for managing large volumes of unstructured data. Additionally, it is media intelligent with numerous features e.g. Avid/Adobe project smarts, camera metadata extraction, AAF/XML/EDL restores and conforms.

Powers Hybrid Workflows for Backup, Disaster Recovery and Archive


With a multitude of data transfer options and flexibility to deploy across on-premise, remote or cloud instances, DNAfabric powers numerous workflows including backup, archive, DR, sharing and collaboration. It powers hybrid workflows spanning multiple locations and cloud infrastructures.

Designed to Simplify Cloud Adoption and Cloud Workflows

DNAfabric is easy to deploy on-premise (physical server, VMware, Virtualbox) or in the cloud (Amazon AMI etc.) enabling users to deploy DNAfabric instances where they need to address data management challenges. It enables mobility of data across local, remote, intra-cloud and inter-cloud pipelines enabling complete flexibility. DNAfabric visibility allows users to keep tight governance of data as it migrates to cloud and remote infrastructures. It further allows them to set and enforce SLAs on their data.


DNAfabric: CLUSTERED DATA SERVICES ARCHITECTURE

METADATA & MANAGEMENT




Centralized Services

DNAfabric is powered by a set of centralized services running on one or more virtual machine instances on-prem or in the cloud. This include a web based console, open xml catalog database, noSQL search DB and an analytics DB.




Single Site Data Managers

Data managers can be deployed on-prem to manage local data management tasks. This include local backup, replication, archive, nearline/project tiering. Additionally, these data managers can scan and index multiple on-prem storage pools for analytics and search.



Multi Site Data Managers

Remote data managers deployed across one or more sites provide high speed synchronization services over UDP. This allows for cross site replication for DR, remote archive and collaboration.



Cloud Data Managers

Cloud data managers provide data management services for cloud pipelines. Visibility services allows for indexing of cloud file-systems and object storage. Data mobility services copies, syncs and moves data between cloud providers or between different classes of storage within a cloud provider.