

The screenshot shows the Komprise website header with navigation links: Products >, Solutions >, Partners >, Resources >, and a green 'Schedule a Demo' button. Below the header is a grid of four team member profiles. Each profile includes a headshot, the person's name and title, and a short bio.

Name	Title	Bio
Kumar K. Goswami	CEO	Kumar has spent 23+ years delivering products that solve complex IT problems with simplicity and cost-efficiency
Krishna Subramanian	President & COO	Krishna is passionate about creating business models that delight and deliver value to customers while being efficient to scale
Michael Peercy	CTO	Mike is a minimalist who takes pride in creating scalable architectures that are easy to manage and just plain work
Mohit Dhawan	VP, Engineering	Mohit brings a proven track record of building engineering teams that deliver killer products

*Kumar Goswami
Analytics-driven data management

BACKGROUND

- Founded 2014
- Platform to manage unstructured data at massive scale
- Large customers manage as much as 100PB
- Customers: large enterprises or companies with lots of data
- Verticals: Financial, Health, Manufacturing, EDA, Entertainment, Government
- IBM and HPE resell Komprise world-wide

6 tenets

- Insight into our data
- Make the insight actionable
- Don't get in front of hot data
- Show us a path to the cloud
- Scale to manage massive quantities of data
- Transparent data movement

3 Architectural pillars

Dynamic data analytics - analyses data so you can make the right decision before buying more storage or backup

Transparent Move Technology - moves data with zero interference to apps, users, or hot data

Direct Data Access - puts you in control of your data - not your vendor

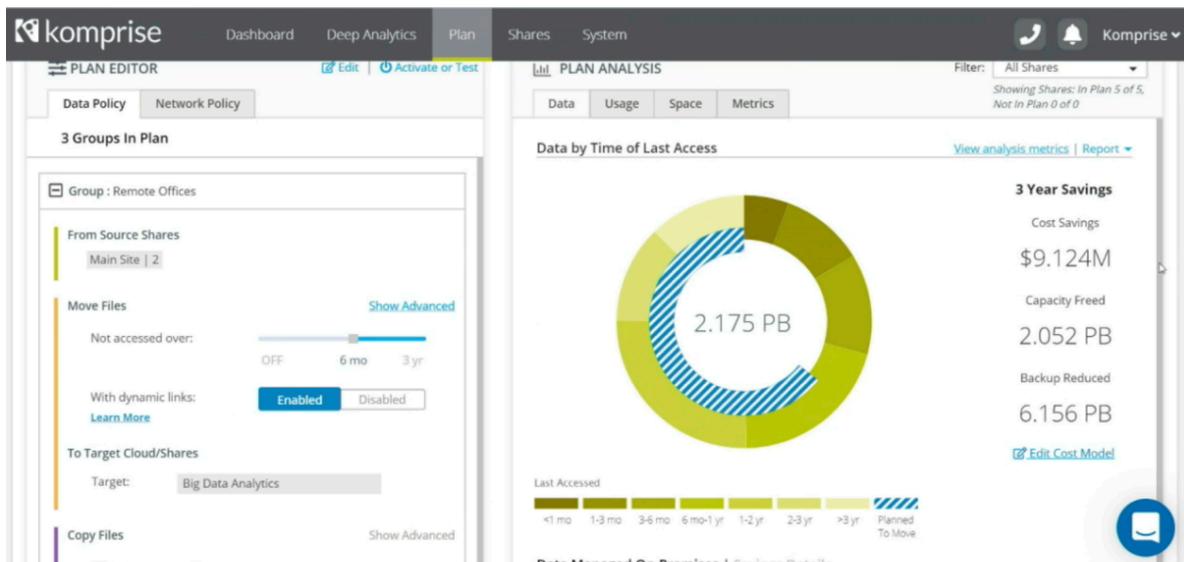
Archive successfully

- no disruption
- transparency
- no interference with hot data
- save money
- without lock-in
- extract value

Managing cloud storage

- IDC - by 2025 50% of all data will be in public clouds
- simplicity of cloud storage exacerbates issues we see on-premises

*Demo



*Mike Pearcy (?), Mohit

Architecture whiteboard

Director - VM can be hosted on-premises or in a cloud
Hosts the console, exposes the API, holds configuration

On-premises is the Observer (can run on ESXi, or can be Linux bare metal)
Discovers the storage (should be in the same DC)

Deep Analytics - indexes the files, Director can run queries against it
Tag the data

Deep Analytics supports multiple Observers (across multiple DCs) - "Global metadata lake"

Automatic performance throttling for scans

“Komprise Dynamic Link” (not a stub) - address hosted by grid of Observers
KCFS - Komprise Cloud File System

Backup? Backup the links, not the data that’s been moved to archive

Largest customer? Close to 100PB

Can have multiple archives

Archive replication? S3 to S3 replication is coming

You can choose to put a second copy somewhere when you’re archiving data

Typical customer starting size is 1PB or more

Cloud?

Spin up Observer instances in the Komprise cloud and point them to Observer instances in the customers cloud

Cloud to cloud migration - e.g. object store to public cloud or back again

“Cloud regret”

One license that includes everything. You can buy a migration-only license separately if you want.