DCIG PULSE

Unstructured Data Management

Komprise Intelligent Data Management

Vendor Strengths Edition

EVALUATION COHORT	12 vendors
ASSESSMENT SCOPE	8 vendor criteria · 68 product capabilities · 44 ecosystem integrations
PUBLICATION	November 2026

Licensed edition prepared for Komprise. Scores and positioning stay fixed at publication and match the market-wide DCIG Pulse.

Market definition

Unstructured data management platforms help organizations understand, organize and better manage data to save costs, modernize storage systems and prepare data for AI and analytics. Managing data includes classifying and curating files, tiering, confining, deleting and migrating files and feeding AI and analytics tools with just the right data. Management platforms sit above the storage layer and work across vendors, sites and clouds.

What the platforms do

These platforms scan file and object storage to build an index of the data. The index shows each file's owner, its age and how often people use it. Some platforms go further and extract and enrich metadata. They scan headers, examine file contents or even run AI to detect sensitive data, capture industry-specific metadata and pull out information for efficient search. Administrators then use that index to find data, tag it and set policies that act on it. Typical actions include tiering cold data to cheaper storage, migrating data between platforms and archiving or deleting data at the end of its life.

Modern platforms must also prepare data for analytics and AI. They select the right files, screen out sensitive content and deliver curated data sets to AI pipelines. The best platforms do this without forcing users to change how they reach their files.

Why buyers adopt them

Unstructured data keeps growing faster than storage budgets, and much of it sits untouched on expensive primary storage. Organizations want to cut storage and backup costs, move data to new platforms or clouds with less risk, and meet retention and privacy rules. They also want to find and use the right data for AI projects.

What DCIG evaluates

DCIG scores 68 product capabilities across the following 8 categories:

- Discovery and Metadata Analytics
- Classification, Tagging and Search
- Policy-Based Tiering and Archiving
- Data Mobility and Replication
- Transparent Access and Namespace
- AI-Ready Data Curation
- Security, Compliance and Audit
- Scale, Deployment and Cost

It also scores 44 third-party integrations across 6 categories, from storage platforms and cloud targets to AI, governance and industry tools.

Inclusion criteria

A platform qualifies for this evaluation when it meets all of the following conditions:

- It discovers and indexes unstructured data across storage systems from more than one vendor.
- It acts on that data through policies, such as tiering, migration, archiving or deletion.
- It supports both on-premises storage and at least one public cloud or object storage target.
- The vendor sells and supports it as a commercial product for enterprise use.

Outside this definition

Storage systems whose data management features work only on their own hardware fall outside this market. So do backup products and standalone migration services that do not index or manage data over time.

Executive summary

The Komprise position

Komprise enters this DCIG Pulse as the leader in vendor-neutral unstructured data management. Its platform indexes data across storage from many vendors, prepares it for AI and moves it without getting in the way of users. Moved files stay reachable at their original paths, with no stubs, agents or rehydration costs.[1–9]

AI sits at the center of that position. Komprise built its AI data services on the same Global Metadatabase that drives tiering and migration. Smart Data Workflows exclude sensitive content, enrich files with extracted tags and deliver curated data to AI pipelines and analytics platforms. KAPPA Data Services and Transparent File Tables run in general availability, and Komprise cites production use in healthcare and higher education.[10–18]

Two design choices set Komprise apart. It extends existing namespaces transparently to secondary storage without fronting data access or imposing a heavy global namespace. It also runs without agents, brittle connectors or customer-run database and search clusters. Lightweight Observers, a global metadata service and one multi-site console do that work instead.[2,6,19–21]

Komprise leads the cohort on five of the 8 vendor criteria: Innovation, Pricing Strategy, Product Capabilities, Business Model and Customer Experience. It ranks second on ISV Ecosystem and Partner Ecosystem, and third on Revenue. Product Capabilities shows its widest lead, at 1.87 points above the cohort mean.

The capability results back that position. Komprise reaches full support on 23 capabilities, the most in the cohort. Komprise builds its platform for enterprise IT deployments at scale, not as a do-it-yourself toolkit for departmental use. It stands alone at full support on six capabilities and ranks first or joint first in six of the 8 capability categories.

FIVE STRENGTHS THAT DEFINE THE KOMPRISE POSITION

1. **Tiering that users never notice.** Patented Transparent Move Technology leaves moved files at their original paths through dynamic links. Users open files as before, and Komprise stays out of the hot data path. Files tiered to read-only targets also sit beyond the reach of ransomware.[1–4]
2. **Data ready for AI.** Smart Data Workflows exclude sensitive content, extract and tag file content and curate the right data for AI pipelines and analytics platforms. Universal File MCP extends the index to AI assistants in early access.[3,10–13,22,23]
3. **One index across every silo, delivered as a service.** Komprise runs the Global Metadatabase as a fully managed, auto-scaled service, with an option to run it in the customer's tenant. Komprise states that hundreds of enterprises use the service, the largest at 100PB, across hundreds of billions of files.[5–9,19,20,24–26]
4. **Migration built for scale and proof.** Parallel data movers, iterative copies and checksum validation on every file support large, low-downtime migrations. Chain-of-custody reports record each file's checksums and timestamps.[5,6,19,20,26,27]
5. **Savings that buyers can measure.** Cost models and showback reports project and track savings by site, share or department. Komprise stands alone at full support on cost savings reporting.[1,6,11,21,28,29]

These strengths suit organizations with large, mixed storage estates that want to cut costs and move data without disrupting users. They also suit teams that need to find and prepare file data for analytics and AI. Komprise also takes a different design path from global file systems and edge caching. The closing section sets out that choice.

DCIG

DCIG PULSE · UNSTRUCTURED DATA MANAGEMENT

Komprise at a glance

How Komprise compares with the 12-vendor cohort

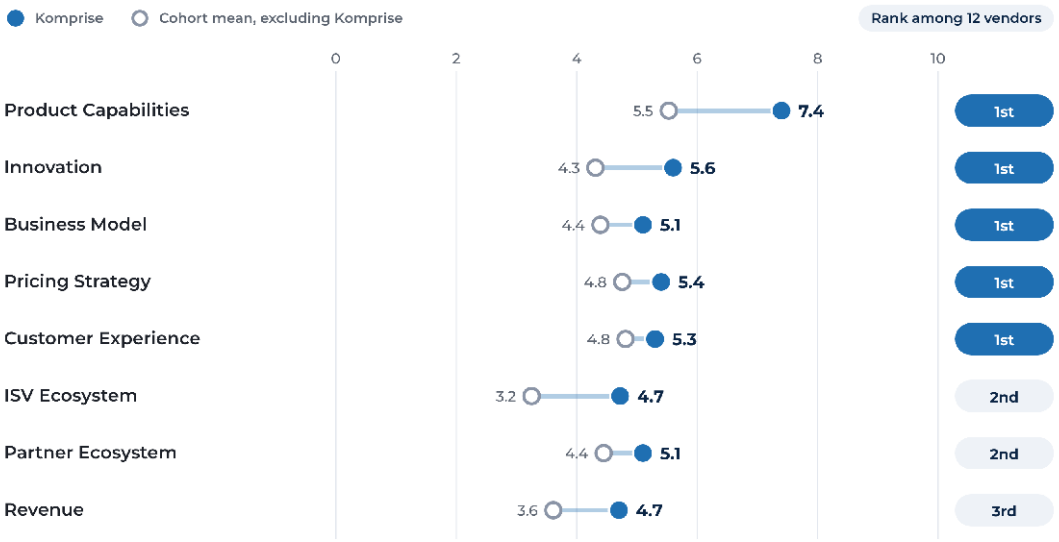
5 / 8
Vendor criteria ranked first
Top 3 in remaining criteria

6 / 8
Capability categories ranked first or joint first
Above the mean in every category

6
Capabilities where only Komprise reaches full support
Including transparent access and path preservation

Vendor criteria

Score on a 0 to 10 scale



DCIG Pulse · Unstructured Data Management · Komprise Vendor Strengths Edition

Source: DCIG evaluation workbook, 12 vendors

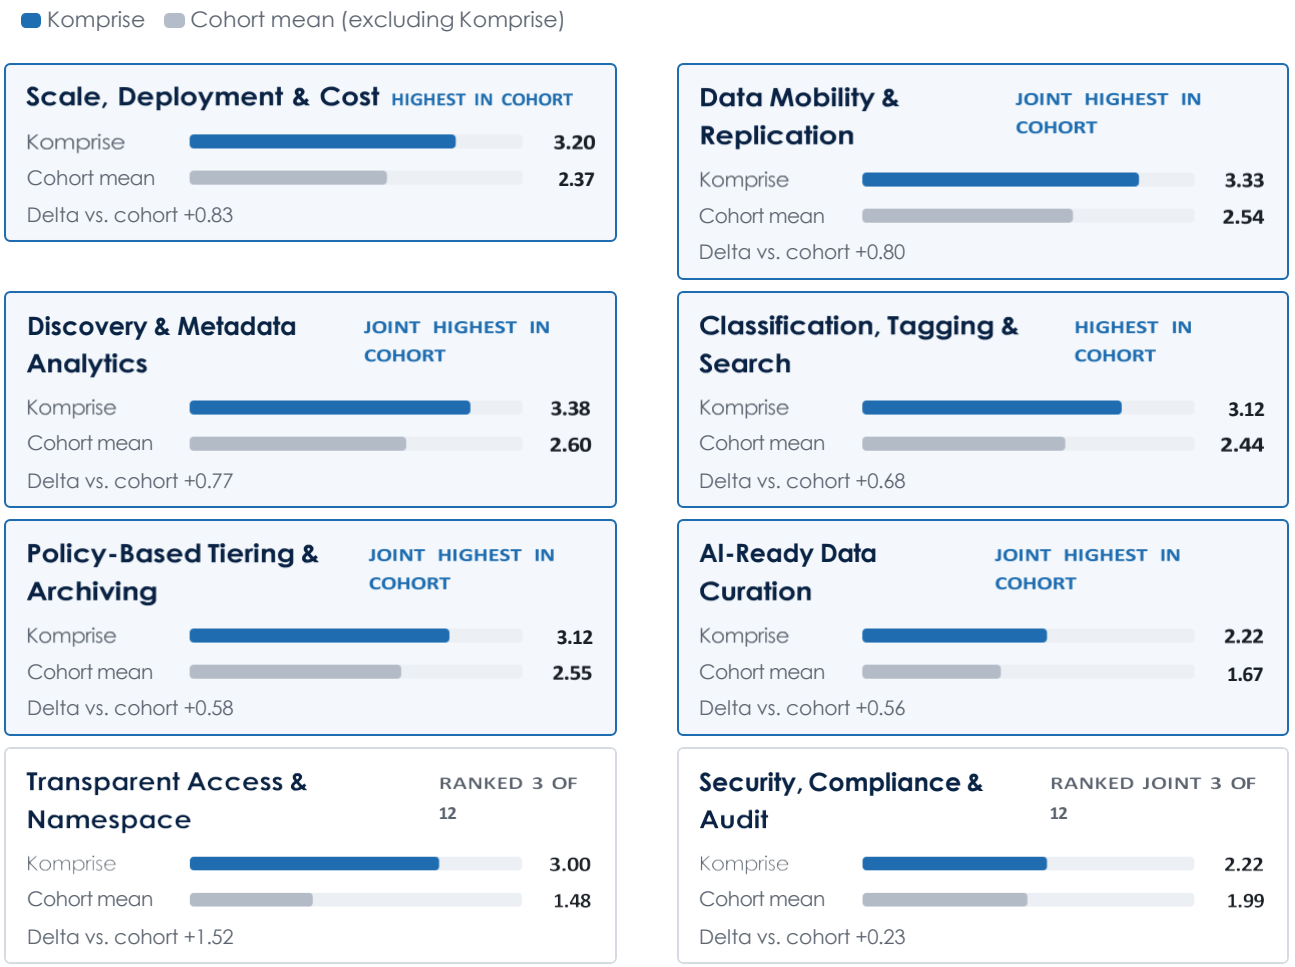
Results by capability category

DCIG groups the 68 product capabilities into 8 categories. Each bar below shows Komprise's average score in a category on the 0 to 4 capability scale. The cohort mean excludes Komprise, and rank shows its position among all 12 vendors.

Komprise ranks first or joint first in six categories, from discovery and classification through data mobility, AI data curation and cost. It leads outright on classification and on cost, and shares first place in the other four. Scale, Deployment and Cost covers architecture, scalability and robustness, and Komprise leads it outright. Komprise sits above the cohort mean in every category, which shows unusual breadth for this market.

Transparent Access and Namespace blends two designs. Komprise stands alone at full support on transparent access to tiered data from the original location and on path preservation. Other vendors score higher on global namespace and multi-protocol access, and some add edge caching. Komprise avoids fronting data access by design, so it scores lower on those capabilities. It never sits in the hot data path, so it never becomes a bottleneck in front of high-performance systems. It also needs none of the heavy infrastructure that global namespaces require.

The transparency of Komprise tiering counts in that category rather than in Policy-Based Tiering and Archiving. Komprise Transparent Move Technology (TMT) tiers data while users and applications keep reaching it from the original location, with native access at every tier. TMT can also archive data to multiple destinations without links.



DCIG

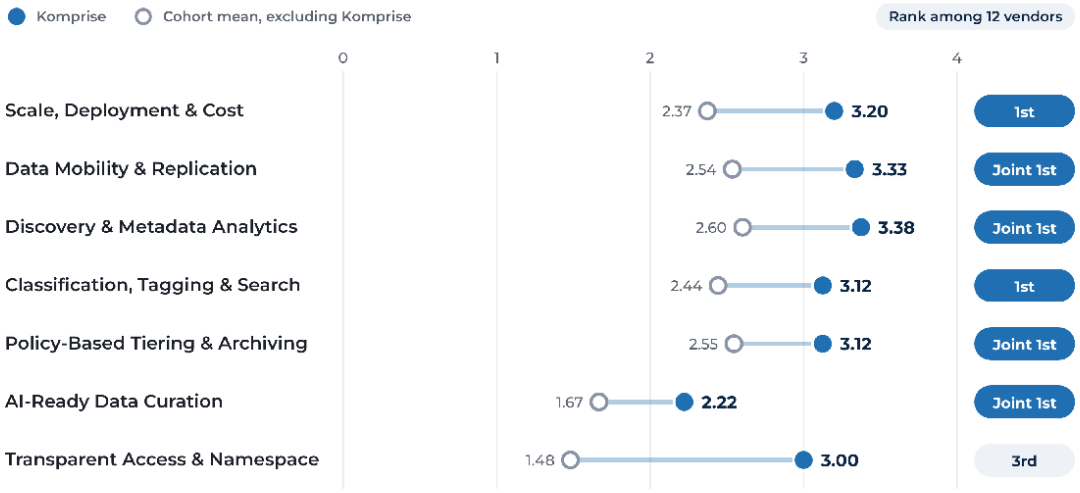
DCIG PULSE · UNSTRUCTURED DATA MANAGEMENT

Where Komprise leads

68 product capabilities scored across 8 categories

Capability categories

Komprise's seven highest-ranked categories · average capability score on a 0 to 4 scale, where 4 denotes full, documented support



Six capabilities only Komprise delivers in full

No other vendor in the 12-vendor cohort reaches full support on these capabilities

Transparent Access
Files stay reachable at their original paths after tiering

Path Preservation
Paths, shares and mount points survive tiering and migration

Multi-Cloud Data Movement
Data moves across on-premises systems and multiple clouds

Curated Dataset Assembly
Right-sized data sets flow to AI and analytics targets

Centralized Management Console
One console spans every site and storage system

Cost Savings Reporting
Savings projected and tracked by site, share or owner

DCIG Pulse · Unstructured Data Management · Komprise Vendor Strengths Edition

Source: DCIG evaluation workbook, 12 vendors

Where Komprise is strongest

The following themes draw on DCIG's analyst commentary for each criterion, capability and integration. Each theme describes what Komprise does differently from its peers and the evidence behind the claim.

1. Tiering that stays out of the user's way

Komprise makes a deliberate design choice: it stays out of the hot data path. Its patented Transparent Move Technology tiers cold files to cheaper storage and leaves a dynamic link at the original location. Users and applications open files exactly as before, and content streams back from the target on access.[1–4,14,30]

By default, a file streams back without returning to primary storage unless someone opens it five times in fifteen days. Administrators can change that rule to suit each workload. Komprise caches recalled data for repeat reads, and it can rebuild links from the object store after a ransomware attack.[1–4,31]

Tiering also protects data. Files tiered to read-only targets sit beyond the reach of ransomware that encrypts primary shares. Komprise states that about 70% of enterprise files have gone cold, so tiering shields the bulk of the estate at low cost.[1,32–34]

Paths, share names and mount points stay valid after tiering. When a customer moves from PowerScale to ONTAP, Komprise converts the links in place. Tiered files keep their paths without rehydration or egress charges. Because no stubs remain, leaving Komprise later carries no re-inflation cost.[1,2,5,6]

Komprise stands alone in the cohort at full support on transparent recall and on path preservation. It shares full support on agentless operation with Hammerspace only, but the two designs differ. Hammerspace places its file system in front of data access through a global namespace. Komprise extends the existing primary namespace to secondary storage and stays out of the data path. That choice avoids the cost and performance overhead of fronting every read. It also matches the cohort high on recall performance.

2. Data prepared for AI

Komprise has extended its index into AI data preparation. Smart Data Workflows select files, run sensitive-data scans, enrich the metadata with content extracted from files and deliver curated sets to AI and analytics applications. Recurring runs rescan only the files that changed since the last run.[10–13,23,35,36]

A healthcare customer offers the strongest proof point. It reduced a 1PB pathology estate to a rolling 33TB working set and removed copies automatically after 30 days. Komprise reports 96% lower cloud cost and ten times faster ingestion for that pipeline.[10–13,17]

Komprise cites further production deployments. Komprise identifies NewYork-Presbyterian as a customer that presented its AI ingestion pipeline at Gartner IOCS. Children's Medical Center Dallas classifies and monitors patient data with Komprise. Duquesne University tags content through an integration with Amazon Rekognition.[17,18,37]

Sensitive data stays out of AI pipelines, because scanners with 68 built-in detectors tag matching files, and workflows exclude or confine them before delivery. Komprise also logs who ingested which data, where and when.[5,9,10,12,14,22,23,33,36,38]

Three 2026 services extend the approach. KAPPA Data Services, now generally available, runs customer-written Python code on each file to extract metadata. Komprise supplies pre-built extraction services, such as DICOM header extraction, for common formats. Transparent File Tables, also generally available, exposes the index to Snowflake and Databricks as Apache Iceberg tables. Iceberg, an open table format, lets analytics engines query the index in place.[2,3,12–16,24,30,36,39]

Universal File MCP, announced on 29 September 2026, lets AI assistants query the index in plain language. It runs in early access, and Komprise targets general availability by the end of 2026. MCP, the Model Context Protocol, gives assistants such as ChatGPT, Claude and Copilot a standard way to reach external data and tools. Every query runs under the user's existing permissions.[2,3,13,14,22,23,30]

Komprise stands alone at full support on curated dataset assembly. It shares first place in the AI-Ready Data Curation category with Hammerspace and Arcitecta. Komprise leads that trio on curated assembly and sensitive content exclusion, while the other two score higher on embedding pipelines and dataset versioning.

3. One index across every storage silo, delivered as a service

The Global Metadatabase sits at the heart of the platform. Komprise delivers it by default as a fully managed, auto-scaled service, with an option to run it in the customer's own tenant. The hosted model spares customers the cost and effort of provisioning, running and scaling an index.[6,9,19,20,24–26]

Komprise states that hundreds of enterprises use the service, the largest at 100PB, with hundreds of billions of files indexed. It also states that its multi-tenant service processes 10PB of data per day per cluster. A grid of virtual appliances called Observers scans NAS, object and cloud storage from many vendors and feeds the index.[5,7–9,19,20,24,27]

Results arrive quickly. A patent-pending sampling technique returns first analytics within about 15 to 30 minutes on petabyte-scale data, while full indexing completes in the background. Adding Observers adds index and scan capacity as the estate grows.[9,19,20,24,27]

Deep Analytics, built on Elasticsearch, lets administrators query every connected system at once by name, path, size, age, owner and tags. Age and access profiles show how much data has gone cold on each share, with age buckets that extend to five or ten years. Universal File MCP adds plain-language queries over the same index.[1,5,6,13,22,24,36,39]

KAPPA extends the index further. Customers apply their own extraction and transformation code to each file as serverless compute, and the results land in the index as searchable metadata.[9,10,14,15,24,36,40]

Komprise reaches full support on billion-file indexing, cross-silo search, access age analytics and multi-vendor scan coverage. It runs billion-file indexing across hundreds of enterprise customers. It matches the cohort high on duplicate detection, orphaned data identification and capacity growth trending. It shares first place in the Discovery and Metadata Analytics category with Starfish Storage. The two match on six of eight capabilities; Komprise scores higher on duplicate detection, and Starfish on direct SQL access to its catalog. Published Starfish material cites

deployments of tens of billions of files, not hundreds of billions. Starfish customers also run its catalog database on their own servers, while Komprise delivers its index as a managed service.

4. Turning queries into actions

Komprise turns a search result into an action without exporting lists between tools. A saved Deep Analytics query can drive a tiering plan, define an AI data set or feed a sensitive-data scan. Smart Data Workflows run that query on a schedule, so Komprise handles new matching data as it arrives. [23,24,35,39]

Policies reach well beyond file age, because a plan can combine owner, department, file type, path, tags and scanner results with date rules. For example, a plan can tier only files tagged to one project that have sat idle for a set period.[5,10,13,24,35,39,41]

The same approach automates classification and enrichment. Teams set up a Smart Data Workflow once, and Komprise keeps discovering, classifying and enriching new files one by one. Policies can then act on the results of earlier analysis, classification and enrichment through saved queries. Patented Elastic Shares technology spreads that work across Observers so every node stays busy. Komprise states that this makes maximum use of available compute, including GPUs.[5,9,10,19,20,23,26,27]

Tagging ties these pieces together, and Komprise applies tags from query results, keyword and pattern rules, the REST API, custom code or third-party AI services. They persist when files move, and they drive tiering, confinement, AI ingestion and reports.[2,9,10,13,14,24,39]

Before anything moves, administrators can model the result. Plan Analysis and the cost comparison tool show capacity, file counts and a three-year cost projection for each scenario. Komprise shares full support on rule-based tagging with Starfish Storage only, and matches the cohort high on policy simulation.[5,21,28,39]

5. Migration built for scale and proof

Komprise Elastic Data Migration runs in repeated passes. An initial bulk copy runs while users keep working. Later delta passes copy only changed data, and a final pass runs after source access stops.[1,5,6,9]

The migration engine scales out, as Observers share work across the grid. An April 2026 patent covers load balancing a single share's work across Observers. A WAN transport called Hypertransfer reduces the effect of latency on long network links. Komprise claims NFS migrations up to 27 times faster than native tools, and SMB migrations up to 25 times faster. Komprise publishes both results in white papers that describe the tests.[5,6,19,20,26,27,42,43]

Proof matters as much as speed, so Komprise validates every file on every pass with MD5 or a hash value. Its chain-of-custody report lists source and destination checksums and timestamps for each file, which suits regulated projects.[5,6,19,26]

Permissions travel with the data, because Komprise preserves NTFS and SMB access lists. It can also translate user and group identities between domains during a consolidation. It stands alone at full support on multi-cloud data movement and shares first place in the Data Mobility and Replication category with HorizonH.[5,6,24,26,44,45]

6. Savings that buyers can measure

Komprise treats cost as a first-class output of the platform. Its cost model converts hardware, backup, software, retention copies, power and staffing into a cost per gigabyte. Plan Analysis then projects savings before any data moves.[1,11,21,28]

Showback reports present total cost, post-tiering cost and savings by site, share or organization. Teams can schedule them as PDF or CSV files for finance and business owners. The Flash Stretch assessment, packaged in March 2026, models how much flash capacity tiering can free up. [5,6,21,28,29,39]

Customer results give the model weight, and Komprise cites Pfizer at 75% lower storage, backup and disaster recovery cost after tiering to Amazon S3. Komprise stands alone at full support on cost savings reporting.[1,11,21,28]

Pricing also stands out for its openness, since AWS Marketplace lists 12-month prices of $61,600 for 250TB and $154,000 for 1PB. Komprise ranks first in the cohort on Pricing Strategy and on the Scale, Deployment and Cost category.[3,25,29,46–48]

7. Light enterprise operations overhead and responsive support

Komprise runs without agents on storage systems, desktops or application servers. It needs no brittle connectors or plug-ins, and IT sets up no database servers, index servers or other complex infrastructure. Komprise uses a hybrid architecture in which Observers deploy as virtual machines on Hyper-V, VMware or KVM, on bare metal or in the cloud. First analysis results appear about 30 minutes after an Observer powers up.[2,6,9,19,20]

The Komprise Director provides one console for every site and storage system. It runs as a cloud service by default or as on-premises software. Komprise also manages the Global Metadatabase, so customers avoid the complexity, overhead and cost of running it. Komprise stands alone at full support on the centralized management console.[6,20,21,25,26]

The grid avoids a single point of failure among Observers. If the Observer managing a share fails, another takes over its requests automatically, and tiered files stay reachable unless every Observer fails. The Director handles configuration and monitoring only, so a Director outage stops neither data transfers nor access to tiered files. Komprise charges nothing per Observer, so customers can add Observers to raise performance. Komprise matches the cohort high on high availability, SaaS delivery and REST API coverage.[2,9,20]

Support terms run clearer than the category norm. Komprise offers 24x7 coverage with a 60-minute response target for critical issues and two hours for high-priority ones. Free assessments show buyers their cold data and projected savings before any commitment. Komprise ranks first in the cohort on Customer Experience.[19,25,28,49–51]

8. A broad storage ecosystem and a focused business

Komprise works with storage through standard NFS, SMB and S3 protocols rather than array plug-ins. That design gives it wide reach across Dell, NetApp, Everpure FlashBlade, Qumulo, VAST Data and Windows file servers. Komprise auto-discovers every share on PowerScale and NetApp clusters during migration setup.[5–9,26,27]

Cloud coverage runs equally deep, with Amazon S3, Glacier, Azure Blob and Azure Files as primary targets. Microsoft lists Komprise as a validated Azure data management partner. Komprise sells

through AWS Marketplace, including AWS GovCloud, and through Microsoft Marketplace. [3,6,9,19,20,25,26,52–58]

Storage vendors also carry Komprise to market. IBM sells it with IBM Cloud Object Storage and Ceph, and HPE resells it with GreenLake and Alletra. Komprise also joined the VAST Cosmos partner community in March 2026. Spectra Logic and Wasabi also appear among its alliance partners. Komprise describes itself as a 100% channel-driven company and ranks second in the cohort on both ISV Ecosystem and Partner Ecosystem.[3,25,52–61]

The business stays focused on this one market. Subscription revenue doubled in each of 2021, 2022 and 2023, and net dollar retention reached 120% for 2022. Deloitte and Inc. growth rankings extend that record into 2025, and the Inc. 5000 listed Komprise for a fifth consecutive year in 2026.[59,62–65]

Product capability differentiators

DCIG scores each of the 68 product capabilities on a 0 to 4 scale, where 4 denotes full, documented support. Komprise reaches full support on 23 capabilities, the highest count in the cohort. The next highest vendor, Starfish Storage, reaches 16. Named Starfish deployments sit mainly in research and HPC environments, where research computing teams rather than central IT typically run it. Komprise focuses on enterprise IT deployments at scale, while letting users talk to and control their own data. The sections below show where Komprise's position stands out most.

● Full support ○ Substantial ○ Partial ○ Limited ○ None

Sole full support in the cohort

Transparent recall and path preservation together carry the patented Transparent Move Technology. No other vendor reaches full support on either one.

● Transparent Access

ONLY VENDOR IN THE COHORT AT FULL SUPPORT

Dynamic links keep tiered files reachable from their original location, and content streams back on access. Customers can configure when files return to primary storage, by default after five opens in fifteen days. Komprise caches recalled data and can rebuild links after a ransomware event. Reads from deep archive classes depend on those classes' retrieval delays.[1–4]

● Path Preservation

ONLY VENDOR IN THE COHORT AT FULL SUPPORT

Komprise can migrate tiered data seamlessly, so tiered files keep their paths. A PowerScale to ONTAP migration, for example, converts link types in place. Object targets preserve the source folder structure as native object prefixes. Komprise respects name and path-length rules on the destination. [1,2,5,6]

● Multi-Cloud Data Movement

ONLY VENDOR IN THE COHORT AT FULL SUPPORT

Komprise moves data between on-premises systems and AWS, Azure, Google Cloud and Wasabi through S3 and file protocols. Observers can run on-premises or in cloud regions, and Hypertransfer reduces latency effects on long links. Buyers still need to plan for egress fees and cold-tier retrieval costs.[6,9,45]

● Curated Dataset Assembly

ONLY VENDOR IN THE COHORT AT FULL SUPPORT

Users build data sets by stacking filters on type, age, owner, tags and scanner results. Scheduled jobs copy only matching files to a chosen target and rescan only changed files. A healthcare customer cut a 1PB estate to a rolling 33TB working set. Komprise claims more than 70% noise removal before ingestion.[10–13]

Centralized Management Console

ONLY VENDOR IN THE COHORT AT FULL SUPPORT

The Komprise Director communicates with the self-monitoring grid, sets policies, monitors activity and reports across every site. Sources from different vendors and locations appear in the same views. The Director runs as a cloud service or as on-premises software. Komprise targets delegated administration for separate business units for early 2027.[6,20,21]

Cost Savings Reporting

ONLY VENDOR IN THE COHORT AT FULL SUPPORT

The cost model computes a per-gigabyte cost from hardware, backup, software, power and staffing inputs. Showback reports present total cost, post-tiering cost and savings by owner. Migration analysis adds a three-year cost projection for each scenario. Cost figures come from models rather than metered consumption.[1,11,21,28]

Rare-air positions

A rare-air position means Komprise holds one of only two full-support scores in the cohort.

Agentless Operation

RARE-AIR POSITION · SHARED ONLY WITH HAMMERSPACE

Komprise needs no agents on storage systems, desktops or application servers. Observers connect over standard NFS, SMB and S3 protocols. Users reach tiered data through links on their existing shares. Unlike the other vendor at full support, Komprise does not front data access.[2,6,19]

Rule-Based Tagging

RARE-AIR POSITION · SHARED ONLY WITH STARFISH STORAGE

Tags come from query results, sensitive-data scanner matches, keyword and pattern rules, the REST API and custom code. Pattern rules can require nearby words, such as a record label near a number. Tags persist when files move and drive tiering, confinement and AI ingestion.[10,13,14,39]

Full support shared with peers

Access Age Analytics	3 of 12 at full support
Cross-Silo Search	3 of 12 at full support
Metadata-Driven Policies	3 of 12 at full support
Minimal-Downtime Cutover	3 of 12 at full support
Multi-Vendor Scan Coverage	3 of 12 at full support
Search-Driven Actions	3 of 12 at full support
Transfer Integrity Validation	3 of 12 at full support
ACL Preservation	4 of 12 at full support
Age & Access Triggers	4 of 12 at full support
Billion-File Indexing	4 of 12 at full support
Object Storage Targets	5 of 12 at full support
Parallel Transfer Engine	5 of 12 at full support
Scale-Out Architecture	5 of 12 at full support
Self-Managed Deployment	5 of 12 at full support
Cross-Vendor Migration	8 of 12 at full support

Joint highest where no vendor reaches full support

No vendor reaches full support on the capabilities below, and Komprise matches the highest score in the cohort on each.

Recall Performance	joint highest · cohort mean 1.3
SaaS Delivery	joint highest · cohort mean 1.3
Bandwidth Throttling	joint highest · cohort mean 1.4
Cross-Domain Identity Mapping	joint highest · cohort mean 1.5
Policy Simulation	joint highest · cohort mean 1.6
Sensitive Content Exclusion	joint highest · cohort mean 1.6
High Availability	joint highest · cohort mean 1.6
Natural-Language Data Search	joint highest · cohort mean 0.8
Content-Based Classification	joint highest · cohort mean 2.0
Capacity Growth Trending	joint highest · cohort mean 2.1
Duplicate Data Detection	joint highest · cohort mean 2.1
Orphaned Data Identification	joint highest · cohort mean 2.1
Administrative Audit Trail	joint highest · cohort mean 2.2
Policy-Aware Selection	joint highest · cohort mean 2.5
REST API Coverage	joint highest · cohort mean 2.6

Ecosystem integration differentiators

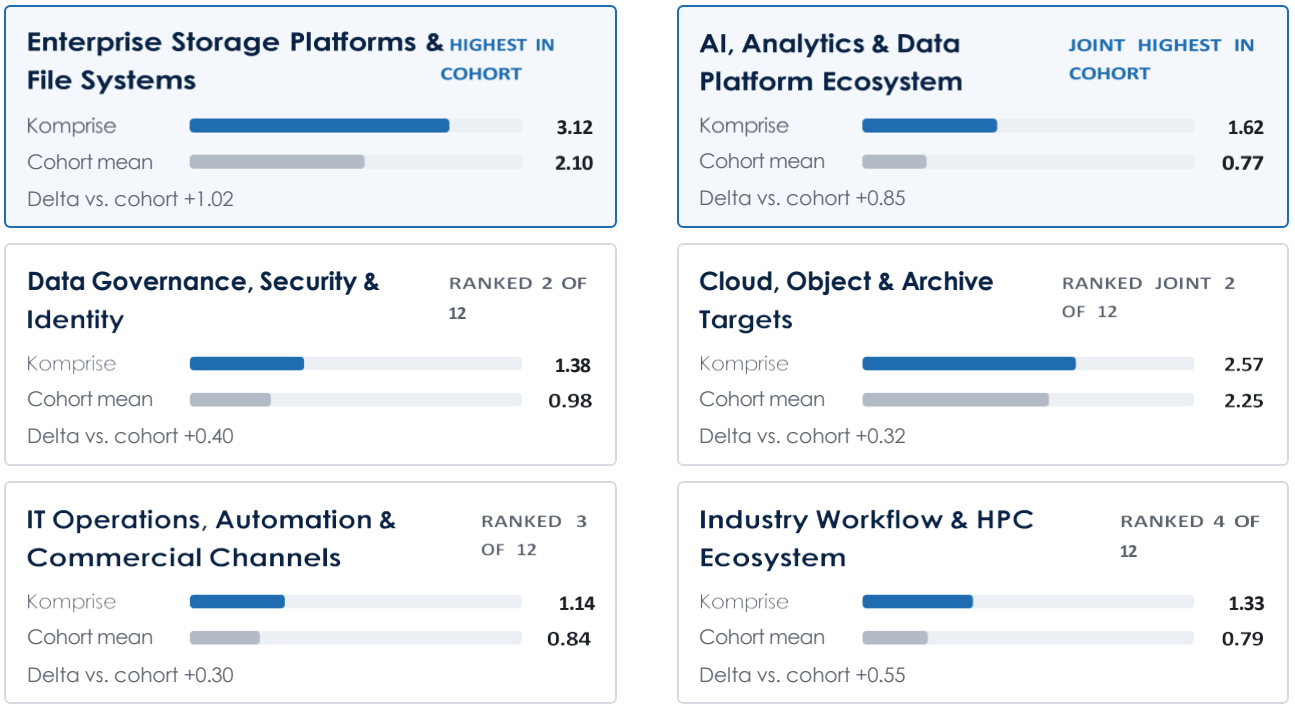
DCIG also scores 44 third-party integrations across 6 categories on the same 0 to 4 scale. Komprise reaches full support on 5 integrations, the highest count in the cohort; no other vendor reaches more than 3. No vendor outscores Komprise on 15 of the 44 integrations.

The pattern shows depth in storage platforms and cloud targets, and a growing reach into AI platforms. The closing section covers the gaps in productivity suites, identity and IT operations tools.

ISV Ecosystem measures technical integration with third-party products, which this section covers. Partner Ecosystem, a separate vendor criterion, measures go-to-market relationships such as alliances, resellers and marketplaces.

Results by integration category

The bars below show Komprise's average integration score in each category on the 0 to 4 scale. Komprise ranks first on enterprise storage platforms. It shares first place with Diskover Data on the AI, analytics and data platform ecosystem.



Leading integrations

Amazon S3 & Glacier

FULL SUPPORT · ONLY VENDOR IN THE COHORT

Amazon S3 serves as a primary Komprise target, with Glacier Deep Archive, Flexible Retrieval and Instant Retrieval supported. Coverage extends to AWS file services such as EFS and FSx, plus AWS GovCloud and Outposts. Administrators choose storage classes at the policy level.[6,9,19,31]

Azure Blob & Azure Files

RARE-AIR POSITION · SHARED ONLY WITH DATA DYNAMICS

Microsoft lists Komprise as a validated data management partner, with guides for migration and tiering to Azure. Migrations can target the Hot, Cool or Archive tier, and tiered data stays reachable through links. Migrations preserve SMB access lists and access times.[5,6,9,19,26,27]

Dell PowerScale & ECS

RARE-AIR POSITION · SHARED ONLY WITH DISKOVER DATA

Komprise supports PowerScale, ObjectScale, ECS, Unity and VNX, with its compatibility matrix revised in September 2026. Adding a PowerScale cluster discovers every access zone, share and export automatically. The integration works at the protocol level rather than through a OneFS plug-in.[5–10,26,38]

VAST Data

RARE-AIR POSITION · SHARED ONLY WITH DATA DYNAMICS

Komprise joined the VAST Cosmos partner community in March 2026. The partnership covers migration into VAST and discovery across silos. It also covers tiering cold data from VAST to archive storage and delivering curated AI data sets. Public detail on supported VAST versions remains limited. [3,25,52–61]

Further integrations where Komprise matches the cohort high

● Windows File Servers & Generic SMB/NFS	joint highest · cohort mean 3.4
○ NetApp ONTAP & StorageGRID	joint highest · cohort mean 2.4
○ Everpure FlashBlade	joint highest · cohort mean 1.6
○ Google Cloud Storage	joint highest · cohort mean 1.8
○ MinIO, Cloudian & Scality	joint highest · cohort mean 2.3
○ AWS, Azure & GCP Marketplaces	joint highest · cohort mean 2.0
○ Okta & SAML/OIDC/SCIM Federation	joint highest · cohort mean 0.6
○ Backup & Cyber-Resilience Platforms	joint highest · cohort mean 0.9
○ AWS Bedrock & SageMaker	joint highest · cohort mean 1.0
○ Google Vertex AI & BigQuery	joint highest · cohort mean 0.3
○ Workflow Orchestrators & CI/CD	joint highest · cohort mean 0.7

Reading the strengths in context

Buyers should weigh the strengths alongside the design choices and boundaries below. DCIG's assessment records the following context for Komprise Intelligent Data Management.

New AI services available for customers

KAPPA Data Services and Transparent File Tables run in general availability for all customers. Universal File MCP runs in early access, with general availability targeted for the end of 2026. Komprise ships no trained classification model of its own and no managed embedding pipeline. Dataset versioning for reproducible training sets also remains limited, and Komprise states that it does not target the smaller model-training market.[2,3,12–14,22,23,30,35,39,66,67]

Security and governance built on a metadata-only design

Komprise stores no customer file data, so data at rest needs no Komprise encryption. The Global Metadatabase keeps metadata encrypted in the Komprise VPC with cloud provider encryption. Data in transit travels encrypted over the internet, and customers can apply their own encryption inside their data centers. Komprise can also encrypt the data it tiers with AES-256.[6,25,26,33]

The platform includes built-in Admin and DA user roles for customer users. Komprise offers no legal hold function, but customers can find data that needs a hold and quarantine it. The multi-tenant Global Metadatabase resides in the United States. Komprise can deploy a customer-specific instance in any AWS region or on premises to meet residency rules.[9,10,13,20,23,25,33,39,41,68]

Komprise holds SOC 2 Type 1 attestation, with Type 2 certification in progress. Komprise ranks joint third of 12 in the Security, Compliance and Audit category, above the cohort mean.

A data manager without a global file system

Komprise built a cross-silo index and chose not to provide one merged file system and namespace. Komprise argues that a global file system places one vendor's file system in front of all storage. In its view, that cancels the performance advantages of premium storage and creates lock-in, as one vendor controls all access. Komprise sets out this position in a published white paper.[2,4,9,20,69]

Replication runs as repeated copy passes with no published recovery point objective. Komprise positions itself as a data management layer, not a disaster recovery product.[5,9,11,26]

Gaps beyond file and object storage

Komprise does not manage content in Microsoft 365, SharePoint or Box, which falls outside its file and object scope. Komprise states that it will support SaaS applications such as Microsoft 365, SharePoint and Box in 2027. Komprise REST APIs expose platform capabilities, so teams can integrate data catalogs, SIEM tools, ServiceNow, Splunk and business analytics systems. Komprise does not support or recommend Terraform or Kubernetes deployments, citing its performance requirements. Directory integration with Active Directory, Entra ID and LDAP ranks sixth, and tape support works only through S3 gateways.[5–9,23,24,26,41]

Commercial considerations

Komprise prices on the customer's total analyzed data footprint. License cost grows as Komprise extends insight and search, including the Global Metadatabase, over more data. The license covers any data management across all analyzed data, so customers keep full flexibility over policies. Komprise states that the cost per terabyte falls as the footprint grows, and it charges nothing per Observer. Overage bills retroactively, and the agreement carries no uptime commitment. The last equity raise dates from January 2023, and Komprise has never disclosed absolute revenue.[3,25,29,46–48,59–65]

Strong independent evidence

In 2026 Komprise joined the Inc. 5000 for the fifth year in a row, after Deloitte Technology Fast 500 listings in 2022, 2023 and 2025. It reached the finals of the Storage Awards for Data Management and Monitoring Vendor of the Year and Storage Optimisation Company of the Year. The Cube Technology Innovation Awards named it a finalist for Most Innovative Data Platform. CRN named it to the 2026 Cloud 100 in the Cloud Storage category.[59,62–65]

Komprise lists alliance partners including AWS, Cloudian, Everpure, Google Cloud, HPE, IBM, Microsoft Azure, NetApp, Nutanix, NVIDIA, Qumulo, Scality, Spectra Logic, SUSE, VAST Data and Wasabi. Its resellers include international partners such as Channel Tools in the United Kingdom and Independent Data Solutions in Australia and New Zealand. Buyers can request reference calls with customers running multi-petabyte estates.[3,52–58,70,71]

Methodology

DCIG Pulse for Unstructured Data Management evaluates 12 platforms. The cohort comprises the following products from the market-wide evaluation:

- Komprise Intelligent Data Management
- Hammerspace Global Data Platform
- Datadobi StorageMAP
- Data Dynamics Zubin
- Arcitecta Mediaflux
- Starfish Storage
- Congruity360 Classify360
- Diskover Data Platform
- Aparavi Data Suite
- Panzura Symphony
- Nasuni DryvIQ
- HorizonH Atempo Miria

DCIG assesses each vendor on 8 vendor criteria, each scored on a 0 to 10 scale. Four criteria cover product strategy and execution: ISV Ecosystem, Innovation, Pricing Strategy and Product Capabilities. Four cover vendor viability: Business Model, Customer Experience, Partner Ecosystem and Revenue.

ISV Ecosystem rolls up technical integrations with third-party products. Partner Ecosystem weighs go-to-market relationships, including alliances, resellers, distributors and cloud marketplaces. Revenue weighs the scale, growth and durability that a vendor's own releases, filings and third-party rankings disclose. Disclosed absolute revenue or recurring revenue carries more weight than growth percentages alone.

DCIG also scores 68 product capabilities across 8 categories and 44 ecosystem integrations across 6 categories. Both use a 0 to 4 scale, where 4 denotes full, documented support. The Product Capabilities and ISV Ecosystem criteria roll up from these detailed scores.

DCIG derives scores from vendor documentation, release notes, press releases and blogs. It also weighs practitioner discussion in community forums where such discussion exists. DCIG reports claims it could not corroborate as vendor statements, and it weighs early-access features below generally available ones.

All comparative figures in this edition come directly from the evaluation workbook. Cohort means exclude Komprise, and rank equals one plus the number of vendors that score strictly higher than Komprise. A joint position means another vendor shares the same score, and a rare-air position means one of only two vendors at full support.

About this edition

This Vendor Strengths Edition forms a licensed extract of the market-wide DCIG Pulse, prepared for Komprise. It replaces the earlier edition and reflects the updated evaluation workbook. This revision also incorporates factual clarifications that Komprise supplied after reviewing the earlier edition. Scores, rankings and commentary stay fixed at publication and match the market-wide report exactly. DCIG altered no score as part of the licensing arrangement, and readers should not treat the analysis as an endorsement of any product.

About DCIG

DCIG, an independent analyst firm, evaluates enterprise storage, data protection, data management and cloud infrastructure technologies. DCIG Pulse reports score vendors and products against published criteria and present findings that IT decision-makers can act on.

DCIG Pulse for Unstructured Data Management
Komprise Intelligent Data Management · Vendor Strengths Edition · November 2026 · Revision 2

© 2026 DCIG. All rights reserved.

Endnotes: sources

Numbers in square brackets throughout this report refer to the sources listed below. Each source comes from the evidence record DCIG kept for the related criterion or capability score, or from material Komprise supplied in its review. Where that record names a source without a web address, the entry appears without a link.

1. Komprise Intelligent Tiering: 3 Steps to ROI
   https://www.komprise.com/blog/komprise-intelligent-tiering-3-steps-to-roi/
2. Transparent Move Technology (Komprise)
   https://www.komprise.com/use-cases/transparent-move-technology/
3. Komprise Universal File MCP
   https://www.komprise.com/mcp
4. Under the Hood: The Komprise Filesystem
   https://www.komprise.com/blog/under-the-hood-the-komprise-filesystem/
5. Komprise Elastic Data Migration: 3 Steps to Lower Risks
   https://www.komprise.com/blog/komprise-elastic-data-migration-3-steps-to-lower-risks/
6. Analyze and migrate your file data to Azure with Komprise Intelligent Data Manager (Microsoft Learn)
   https://learn.microsoft.com/en-us/azure/storage/solution-integration/validated-partners/data-management/komprise-quick-start-guide
7. Komprise Compatibility (modified 26 Sep 2026)
   https://www.komprise.com/compatibility/
8. Komprise Integrations (modified 2 Sep 2026)
   https://www.komprise.com/integrations/
9. Komprise Technical Overview
   https://www.komprise.com/resources/komprise-technical-overview-read-online/
10. Komprise Smart Data Workflows: 3 Steps to AI-Ready Data
    https://www.komprise.com/blog/komprise-smart-data-workflows-3-steps-to-ai-ready-data/
11. Unstructured Data Quality's Role in AI Inferencing Costs
    https://www.komprise.com/blog/unstructured-data-qualitys-role-in-ai-inferencing-costs/
12. Komprise Intelligent AI Ingest
    https://www.komprise.com/use-cases/ai-data-ingestion/
13. Ask Your Storage a Question: 18 Prompts for IT with Komprise Universal File MCP
    https://www.komprise.com/blog/ask-your-storage-a-question-18-prompts-for-it-with-komprise-universal-file-mcp/
14. Komprise KAPPA Data Services
    https://www.komprise.com/product/kappa-data-services/
15. Komprise AI Preparation and Process Automation for DICOM Files
    https://www.komprise.com/blog/komprise-ai-preparation-and-process-automation-for-dicom-files/
16. Komprise AI Preparation and Process Automation for Image Files
    https://www.komprise.com/blog/komprise-ai-preparation-and-process-automation-for-image-files/
17. Komprise blog: Healthcare AI Data Pipelines Explored at Gartner IOCS, 13 Jan 2026
    https://www.komprise.com/blog/healthcare-ai-data-pipelines-explored-at-gartner-iocs/
18. Komprise case study: Duquesne University, AI tagging with Amazon Rekognition, 2024
    https://www.komprise.com/wp-content/uploads/komprise-duquesne-case-study-2024.pdf
19. Komprise FAQs
    https://www.komprise.com/resources/faqs/
20. Closer Look: Komprise Elastic Grid
    https://www.komprise.com/blog/closer-look-komprise-elastic-grid/
21. Komprise Reports: Turn Storage Data Into FinOps Evidence
    https://www.komprise.com/blog/komprise-reports-3-steps-to-cost-saving-insights/
22. Komprise Breaks the AI Context Barrier and Reduces MCP Bloat with New Universal File MCP Solution (29 Sep 2026)
    https://www.komprise.com/komprise-breaks-the-ai-context-barrier-and-reduces-mcp-bloat-with-new-universal-file-mcp-solution/
23. Interview: Komprise Universal File MCP
    https://www.komprise.com/blog/interview-komprise-universal-file-mcp/

24. Komprise Deep Analytics
https://www.komprise.com/resources/komprise-deep-analytics-read-online/
25. Komprise Business Agreement - SaaS (last updated June 2026)
https://www.komprise.com/business-agreement/
26. AWS Prescriptive Guidance: Komprise Elastic Data Migration (partner-supplied attributes, last updated 15 Nov 2024)
https://docs.aws.amazon.com/prescriptive-guidance/latest/migration-tools/data-mobility-kompise.html
27. Komprise Awarded Elastic Shares Patent (29 Apr 2026)
https://www.komprise.com/blog/new-komprise-patent-solves-the-idle-compute-problem-in-unstructured-data-processing/
28. New Komprise Flash Stretch Enables Enterprises to Stretch Existing Hot Storage Capacity (26 Mar 2026)
https://www.komprise.com/press-releases/new-komprise-flash-stretch-enables-enterprises-to-stretch-existing-hot-storage-capacity/
29. Komprise Intelligent Data Management
https://www.komprise.com/product/intelligent-data-management/
30. Komprise press releases: 'Komprise Accelerates Agentic AI with Serverless Compute for Unstructured Data' (19 Feb 2026), Elastic Shares patent (29 Apr 2026), 'Komprise Delivers Query-Ready Enterprise Unstructured Data to AI and Data Lakehouses' (23 Jun 2026), 'Komprise Breaks the AI Context Barrier… Universal File MCP' (29 Sep 2026), VP Product Management hire (13 Oct 2026)
31. Komprise and Spectra Logic: Archiving Data with Insight
https://www.komprise.com/wp-content/uploads/2018/03/4952.pdf
32. Object Lock: Immutable Storage for Data Protection (Komprise)
https://www.komprise.com/glossary_terms/object-lock/
33. Komprise Data Security, Compliance and Governance
https://www.komprise.com/product/data-security-compliance/
34. The Komprise Data Experience (KDX)
https://www.komprise.com/kdx-komprise-data-experience/
35. Komprise Smart Data Workflows
https://www.komprise.com/product/smart-data-workflows/
36. Komprise Transparent File Tables
https://www.komprise.com/product/transparent-file-tables/
37. Komprise customer story: Children's Medical Center Dallas, visibility and classification for ePHI data
https://www.komprise.com/resource/healthcare-medical-center-customer/
38. Using Regex to Improve Data Search and Classification with Komprise
https://www.komprise.com/blog/using-regex-to-improve-data-search-data-classification-with-komprise/
39. Search and Act on Unstructured Data With Deep Analytics
https://www.komprise.com/blog/komprise-deep-analytics-3-steps-to-actionable-queries/
40. Komprise AI Preparation and Process Automation for Genomics Files
https://www.komprise.com/blog/komprise-ai-preparation-and-process-automation-for-genomics-files/
41. Komprise documentation portal (login required)
https://docs.komprise.com/
42. Komprise white paper: Elastic Data Migration for NAS, NFS migration performance
https://www.komprise.com/wp-content/uploads/Elastic_DM_NASmigration.pdf
43. Komprise white paper: Hypertransfer for SMB migration performance
https://www.komprise.com/wp-content/uploads/Komprise-Hypertransfer-Migration-White-Paper.pdf
44. Komprise adds automated share and permission mapping to speed data migrations (Blocks and Files)
https://blocksandfiles.com/2025/03/25/komprise-data-migration/
45. Komprise Elastic Data Migration
https://www.komprise.com/product/elastic-data-migration/
46. Microsoft Learn "Azure Storage Migration Program" partner terms (updated 22 Apr 2026)
47. komprise.com/azure-migration program page
https://www.komprise.com/azure-migration
48. "New Komprise Intelligent Tiering for Azure… Exclusive Microsoft Azure Marketplace Offer" (27 Apr 2023)
49. Komprise support page and Freshdesk portal
https://www.komprise.com/support/
50. Komprise Technical Professional training
https://www.komprise.com/training/

51. 'Office Hours' blog series (Jul-Sep 2026)
https://www.komprise.com/blog/
52. komprise.com/partners (Nov 2025) and komprise.com/partnerprogram (Aug 2025)
https://www.komprise.com/partners
53. Komprise blog, 'The Komprise Konnect Partner Program in the Spotlight' (11 Nov 2020)
54. Komprise blog, 'Komprise and HPE GreenLake' (updated 19 Dec 2025)
55. IBM Komprise product page and IBM Cloud catalog
56. AWS Marketplace and Microsoft Marketplace listings
57. CRN Cloud 100 (13 Jan 2026)
58. VAST Cosmos partnership (13 Mar 2026)
59. Komprise Raises $37M (24 Jan 2023)
60. Komprise press releases page: VAST Cosmos partnership (13 Mar 2026), serverless compute for agentic AI (19 Feb 2026), Transparent File Tables (23 Jun 2026), Universal File MCP (29 Sep 2026), VP Engineering hire (15 Jan 2026), VP Product Management hire (13 Oct 2026)
61. Komprise about-us leadership page
62. Komprise Doubles Revenues in 2022, 24 Feb 2023
63. Komprise Doubles Subscriptions Again in 2023, 14 Feb 2024
64. Deloitte Technology Fast 500 release, 19 Nov 2025, and Deloitte 2025 rankings list
65. Inc. 5000 releases, 13 Aug 2024, 12 Aug 2025 and 12 Aug 2026, and Inc.com company profile (ranks and three-year growth)
66. Komprise Joins Vast Cosmos (StorageNewsletter)
https://www.storagenewsletter.com/2026/03/23/komprise-joins-vast-cosmos-as-a-technology-partner-for-unstructured-data-classification-lifecycle-management-and-ai-ingestion/
67. Komprise Transparent File Tables announcement (StorageNewsletter)
https://www.storagenewsletter.com/2026/06/24/komprise-delivers-query-ready-enterprise-unstructured-data-to-ai-and-data-lakehouses-without-moving-a-single-file/
68. Komprise for Legal Services
https://www.komprise.com/industries/legal-services/
69. Komprise white paper: Global Namespace vs Global File System
https://www.komprise.com/resource/namespace-global-file-system/
70. Channel Tools (UK): Komprise portfolio page
https://channel-tools.co.uk/portfolio/komprise/
71. Independent Data Solutions (ANZ): Komprise products and services page
https://ids-g.com/products-services/komprise/