

Lightning-fast File Operations for Extremely Large Namespaces

Techniques for Applying Structure to Unstructured Data

May 17th, 2017

Mitch Crane, VP of Product Engineering at Cloudfenna, will illustrate the power of separating content and metadata. When metadata is extracted into a database, the index powers web-scale file access, search, and audit in ways not before possible.



We're dealing with massive amounts of
Unstructured Data

Segmented namespaces



Single namespace

Virtual namespace via metadata



Creating a single namespace

1

Normalize
Unstructured Data

2

Centralize Metadata
Store

Separate content and metadata
(inode in the cloud)



Lightweight
Metadata

Creating a single namespace

1

Normalize
Unstructured Data

2

Centralize Metadata
Store



- ✓ Single database for multiple data silos
- ✓ Massively scalable
- ✓ High performance and responsive

1 File Access



Anywhere, any device access to on-prem and cloud file repositories

2 File Search



10x faster file sharing

3 Audit & Governance



Share files between data silos:
Microsoft Office Online


1 File Access

2 File Search

3 Audit & Governance



Search multiple data repositories with one query




Machine learning on the Cassandra database makes search smarter

1 File Access

2 File Search

3 Audit & Governance



Take your file analysis and data security to the next level. Recognize user behavior trends to alert you about security and compliance risks.

- Data classification
- User behavior triggers
- Security and compliance checks

Lightning-fast File Operations for Extremely Large Names Spaces

Techniques for Applying Structure to Unstructured Data

May 17th, 2017

Mitch Crane, VP of Product Engineering at Cloudfenna, will illustrate the power of separating content and metadata. When metadata is extracted into a database, the index powers web-scale file access, search, and audit in ways not before possible.