

adaptTo()

APACHE SLING & FRIENDS TECH MEETUP

BERLIN, 26-28 SEPTEMBER 2012

Jackrabbit Oak

the next generation content repository

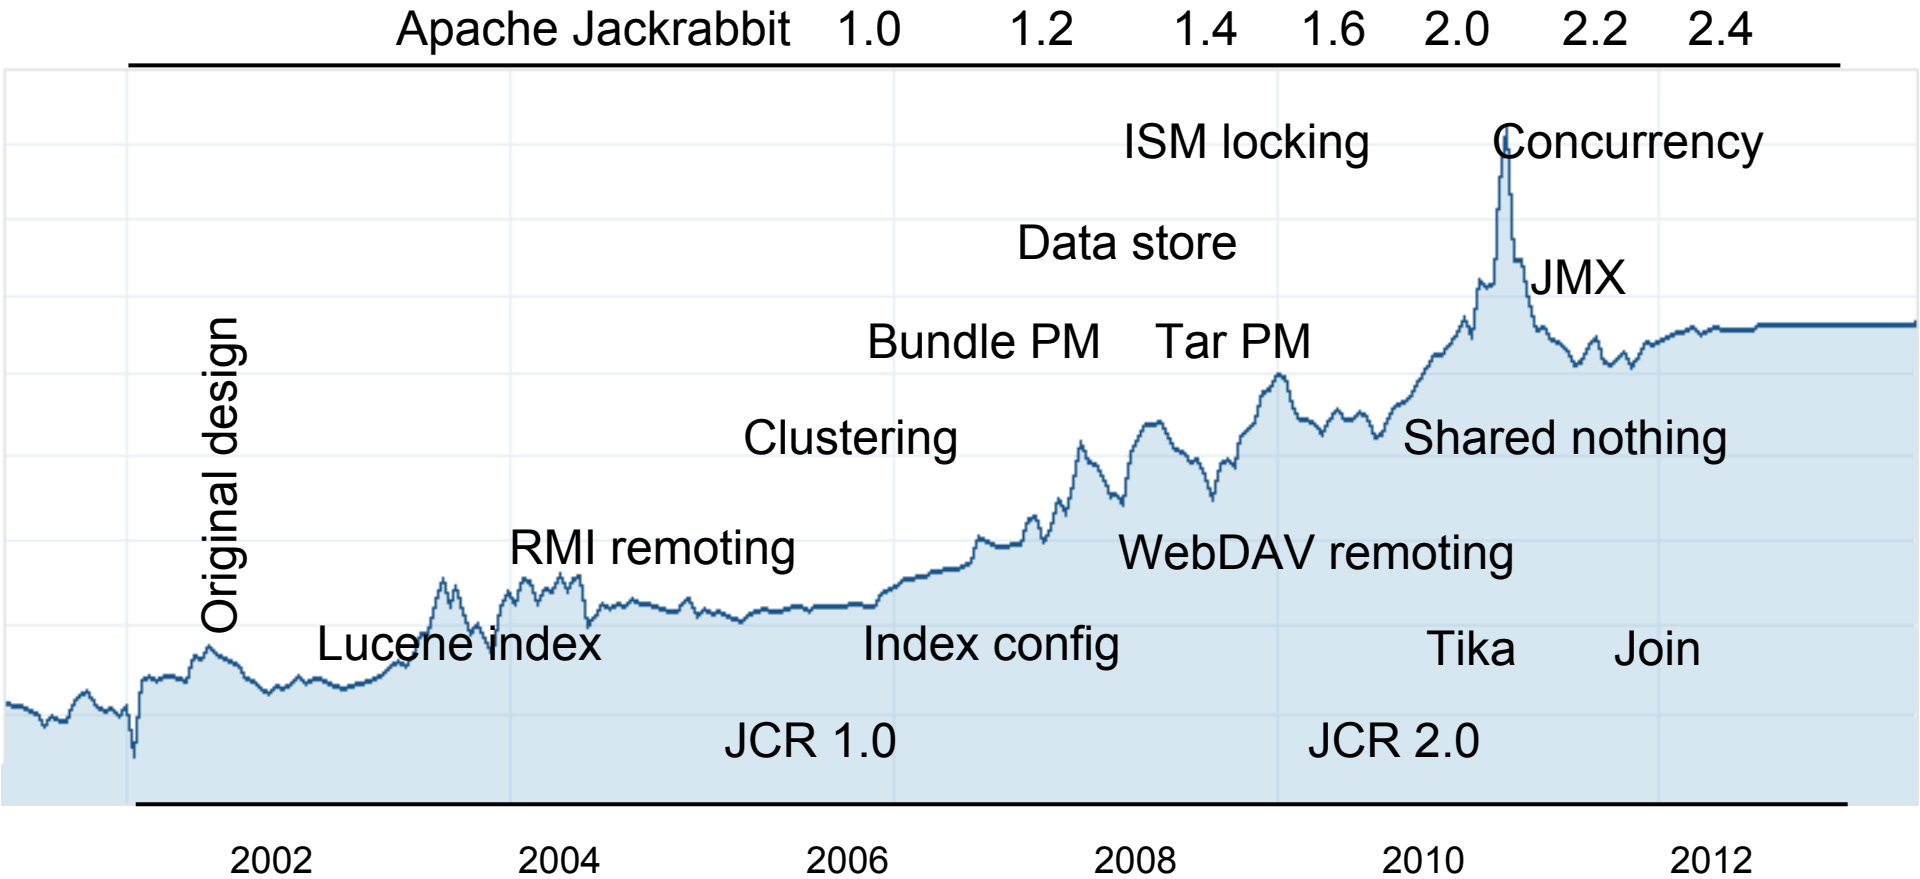
Jukka Zitting, Michael Dürig
{jzitting, mduerig}@adobe.com



Adobe



Jackrabbit feature timeline



Design trajectory



Time for a redesign



Time for a redesign



Time for a redesign



- **Brief history**
 - Early ideas discusses already 2008
 - Prototyping work in 2011
 - Project Oak started early 2012
 - Currently 2 - 8 people actively working on Oak
 - up from 2-3 in early Spring

■ Key goals

- Backwards compatibility (to some degree)
- Performance especially for concurrent access
- Scalability for huge repositories (> 100M nodes)
 - Large number of child nodes (flat hierarchies)
- Support managed environments (OSGi, Spring)
- Cloud deployments

■ Non-goals

- Hard consistency guaranteed
- Full JCR coverage

■ Subproject of Apache Jackrabbit


<http://jackrabbit.apache.org/oak/>

oak-dev@jackrabbit.apache.org

<http://svn.apache.org/repos/asf/jackrabbit/oak/>

<https://github.com/apache/jackrabbit-oak>

Technical overview



JCR

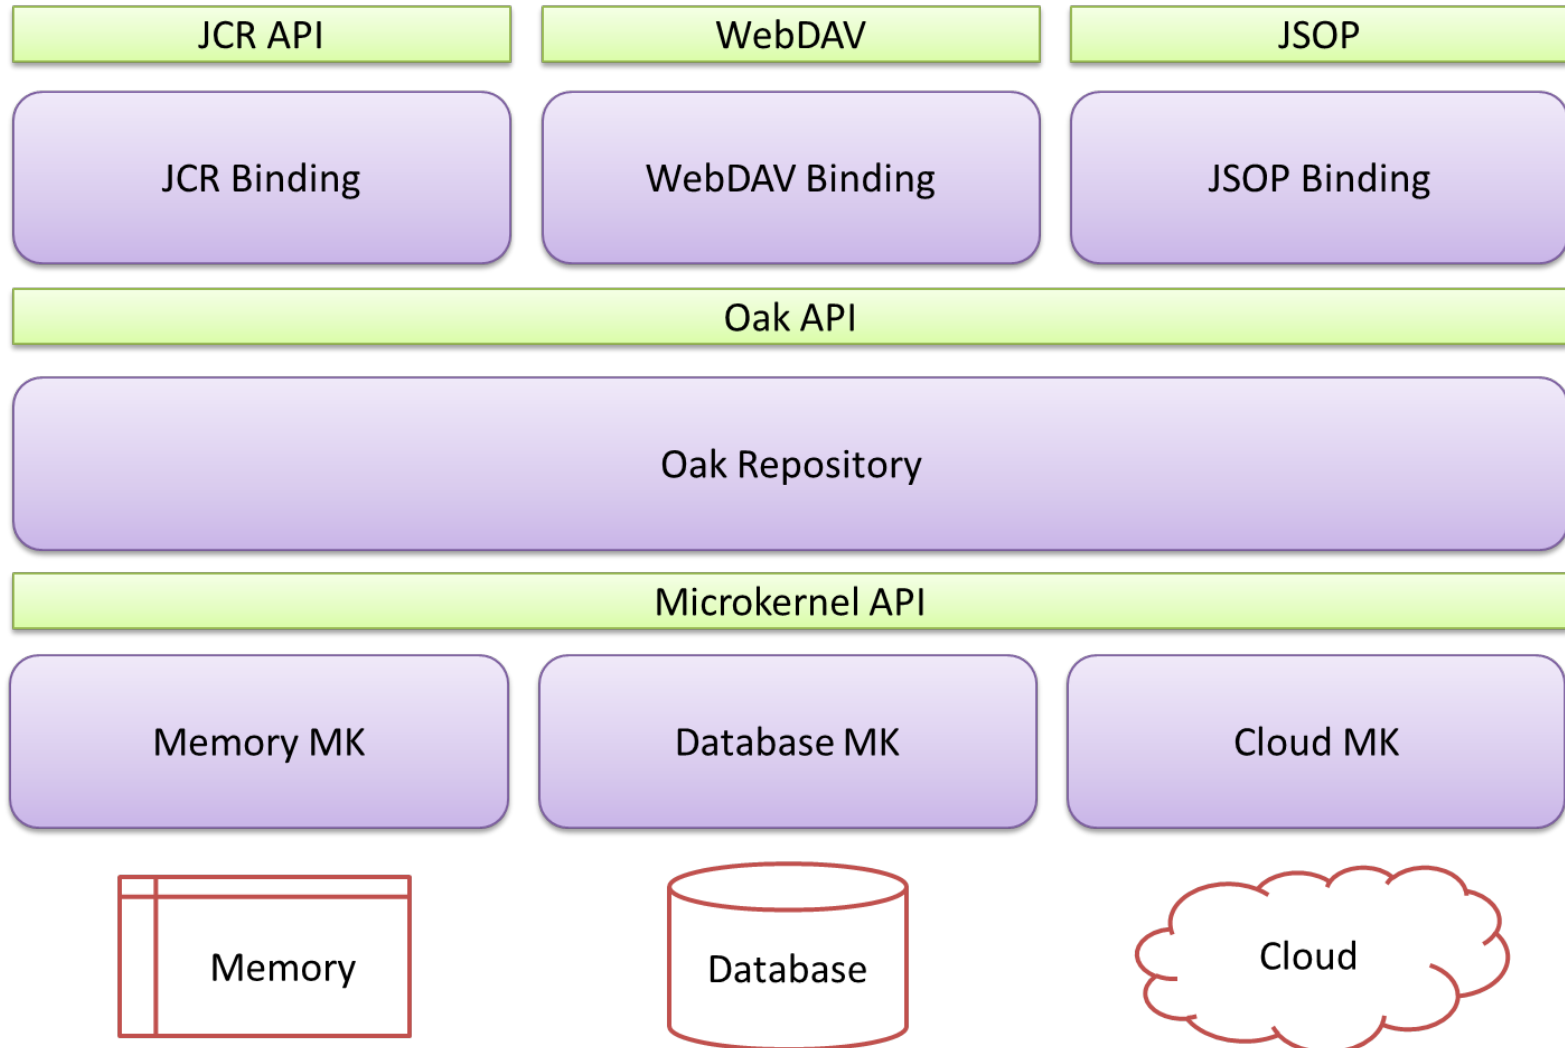
WebDAV

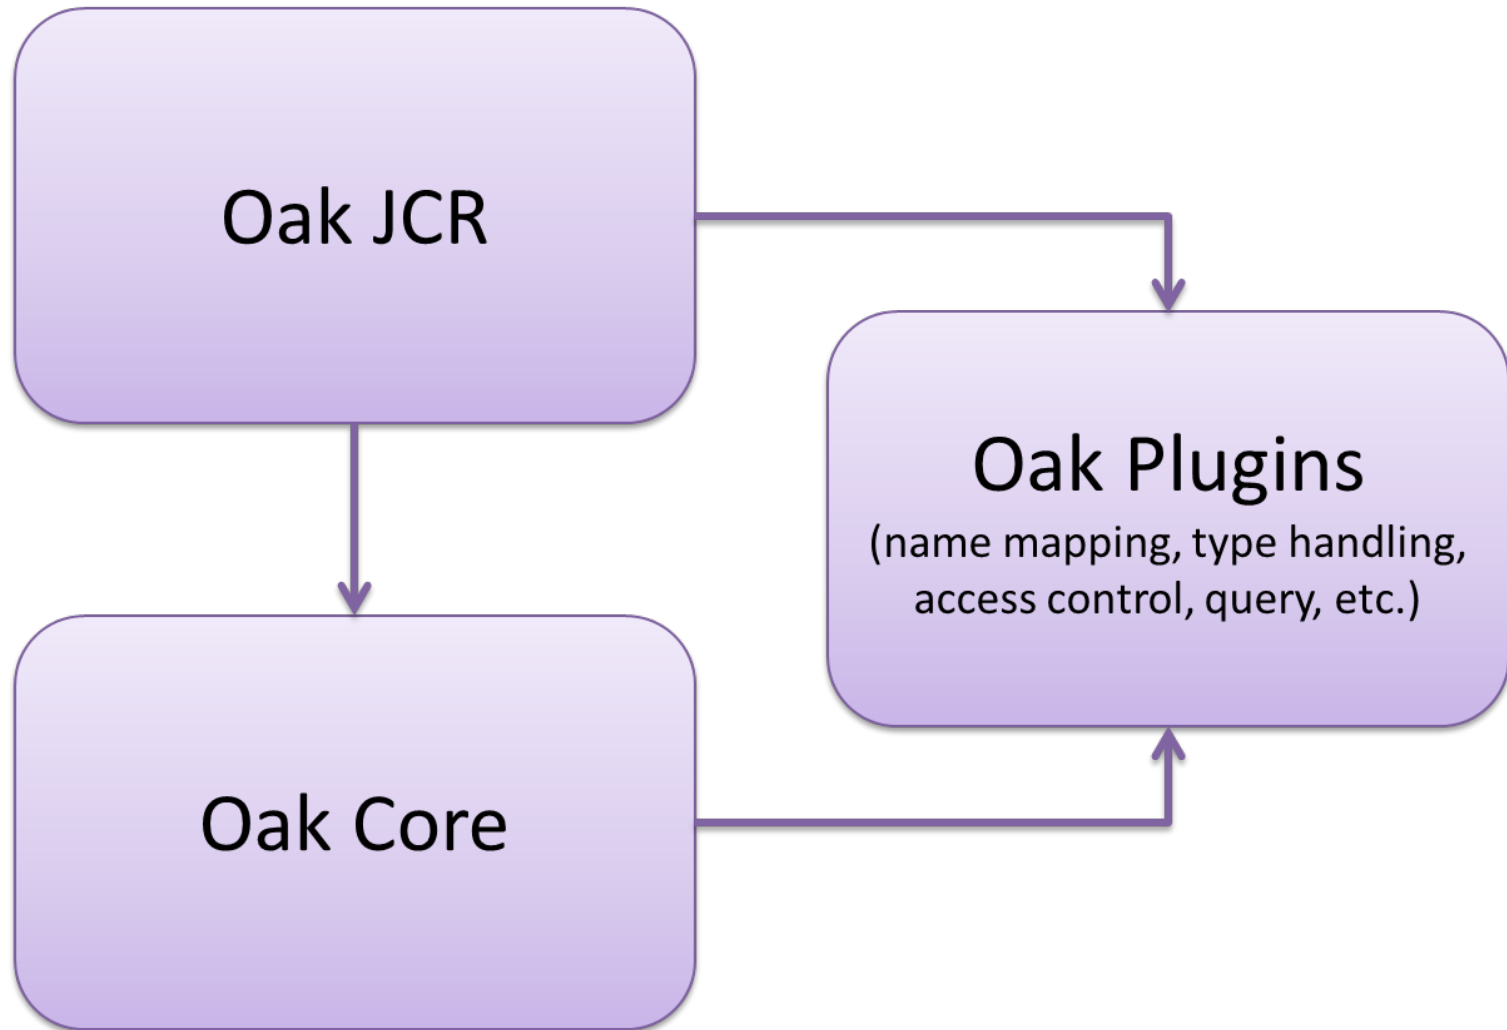
JSOP

Oak

Storage layer (native FS, NoSQL, JDBC, etc.)

Technical overview





- Lightweight tree model
 - Plugins for everything else
- MVCC
 - Stable, immutable snapshot for sessions
 - Snapshot isolation
 - relaxed first committer wins strategy [1]
- Branch and merge
 - Support for large transactions
 - Persisted transient space

[1] [http://wiki.apache.org/jackrabbit/Transactional model of the Microkernel based Jackrabbit prototype](http://wiki.apache.org/jackrabbit/Transactional%20model%20of%20the%20Microkernel%20based%20Jackrabbit%20prototype)

- **Microkernel [1]**
 - Think hierarchy aware persistence manager
 - Stateless
 - Scalable
 - JSON based API
 - Remotable
 - Interchangeable
 - Portable

[1] <http://wiki.apache.org/jackrabbit/RepositoryMicroKernel>

Key concepts

- **Node state model [1]**
 - Low level abstraction of tree content
 - Immutable
 - Thread safe
- **Tree model [2]**
 - Mutable tree model
 - Builders for node states
 - Transient space

[1] <https://github.com/apache/jackrabbit-oak/blob/trunk/doc/nodestate.md>

[2] <https://github.com/apache/jackrabbit-oak/blob/trunk/oak-core/README.md>

- **Plugins**
 - **Commit hooks**
 - Validators
 - Editors
 - **Indexing**
 - **Access control, permissions, privileges**
 - **User management**
 - **Namespaces**
 - **Conflict resolution**
 - **Uniqueness constraints**
 - **Node type validation**

Trade offs



Trade offs

- Trade consistency for availability
 - Write skew [1]
 - Cluster sync
- Limited support for SNS
 - Auto created node names
 - Naming convention
- Limited support for orderability
 - Graceful degradation

[1] [http://wiki.apache.org/jackrabbit/Transactional model of the Microkernel based Jackrabbit prototype](http://wiki.apache.org/jackrabbit/Transactional%20model%20of%20the%20Microkernel%20based%20Jackrabbit%20prototype)

- **Observation**
 - **Affected by session isolation**
 - auto-refresh when observation is used
 - **Effect based**
 - exact sequence of operations not reflected
 - **Limited support for session specific information**

Where we are now



Current status

- **What works already?**
 - Sling core components
 - Granite/CRX quickstart
 - CRXDE Lite
 - CRX package manager

- **What are we working on?**
 - CQ5 quickstart
 - Geometrixx
- **What's next?**
 - More features: security, versioning, etc...
 - Performance/scalability
 - Testing and optimization
 - Cloud deployment and clustering

Getting involved



How can you help?

- Report on success on failure
 - Is the stuff you're working on available on Oak?
 - Does it work as expected?
 - Any differences to Jackrabbit 2?
- Describe benchmarks
 - performance, scalability
- Tell us about useful new extension points

Further information

- Website: <http://jackrabbit.apache.org/oak>
- Source code:
<https://github.com/apache/jackrabbit-oak>
- Mailing lists:
{oak-dev, oak-issues, oak-commits}
@jackrabbit.apache.org
- Wiki:
<http://wiki.apache.org/jackrabbit/Jackrabbit 3>

OAK

Jackrabbit Oak – the next generation content repository.



Jackrabbit Oak is an effort to implement a scalable and performant hierarchical content repository for use as the foundation of modern world-class web sites and other demanding content applications. The Oak effort is a part of the **Apache Jackrabbit** project. Jackrabbit is a project of the **Apache Software Foundation**. Oak is currently alpha-level software. Use at your own risk.

DOWNLOAD

The latest Oak sources are available for checkout from **svn**, or you can **fork them** on GitHub.

See also our monthly releases on the Jackrabbit **download page** for slightly more stable versions of the codebase.

LEARN

Many parts of Oak are still under construction, so it may be a bit difficult to find your way around the codebase. The **README files**, the **Jackrabbit 3 wiki page**, and the Oak **mailing list archives** are good places to start learning about Oak.

PARTICIPATE

The best place for Oak-related discussions is the **oak-dev@** mailing list. To subscribe, send a message to **oak-dev-subscribe@**.

Use the **OAK issue tracker** to submit issues, comments or patches.