

2. Tickets resolved this week:

key	summary	type	created	updated	due	assignee	reporter	priority	status	resolution
Unable to locate Jira server for this macro. It may be due to Application Link configuration.										

3. Tickets created this week:

key	summary	type	created	updated	due	assignee	reporter	priority	status	resolution
Unable to locate Jira server for this macro. It may be due to Application Link configuration.										

Minutes

Child Resources

- Assumption: When you make a request on a container you should know how many children that container has
- Fedora should provide this information
- Going into a paging scenario with this information would be nice
- The current code iterates over all the children and counts them
 - This can be a performance issue with very large numbers (millions) of children
- Instead of dynamically creating this count, we could maintain a property that gets updated when children are created or removed
 - Right now, we don't have good information on when children are created via federation
 - In general, it is difficult to know how many children are in a federated file system period. This should not stop us from making in-repository storage perform better
- There is a child count property we could use - it does not indicate whether those children are in repository or federated storage
- Could we use a listener to listen for federated changes and update the property in Modeshape?
 - Record number of children, bytes, and a timestamp and us a TTL to update at periodic intervals
 - Not sure if we could override the way the federated filesystem gets queried for its contents
 - We could pursue one solution for in-repository content and another solution for federated resources
 - Both solutions should use and maintain the same property
- **Action:** Create a separate ticket for each solution

ModeShape 5 Update

- There will probably be a couple more 4.x releases
- ModeShape 5 is projected to be released in March
- We will want to install and test this release with Fedora 4 when it is available

Fedora Specification

- A message regarding the [API spec](#) was sent to the Fedora community yesterday
- We want to clearly define what Fedora is in a specification
- The proposal includes API and SPI services and WebAC for authorization
- An effort is currently underway to detail each of the services
- We want to have a number of alternate implementations to verify that the API is independent of any particular implementation
- We need to have TCKs that progress alongside the specification work
- Who will volunteer to take responsibility for the service pages on the wiki?
 - [Resource Versioning](#): A. Soroka
 - [Resource CRUD](#): Benjamin Armitor
 - [Atomicity](#): Jared Whiklo
 - [Authorization](#): Allen Flynn, Namita
 - [Binary Fixity](#): Andy Wagner
 - [Atomic Batch Operations](#): Jared Whiklo
 - [Messaging SPI](#): Unknown User (acoburn)
- What will the final spec look like and where will it live?
 - It should be in the same form as the other W3C recommendation documents that are out there
 - It should not just live in GitHub or on a wiki

- W3C Respec for authoring: <https://www.w3.org/respec/>
- This is difficult to do in the wiki. Probably still difficult in GitHub
 - Maybe a Google Doc?
 - Useful but will not help produce a spec doc
 - Are GitHub issues a barrier?
 - Could we use JIRA instead since we already have a running instance?
- **Action:** As a starting point we will use a Google Doc but transfer to something else if need be

MySQL Migration Testing

- There are some documented issues with LevelDB
- Need to migrate backup dump from LevelDB to few Fedora repo v
- Application Link configuration.

ation from LevelDB to MySQL

Unable to locate Jira server for this macro. It may be due to

Fedora Backup

- There is no guarantee of order of resources in a Fedora backup dump