

# 2014-08-13 - Fedora Technical WG Meeting

## Time/Place

- Time: [12:00pm Eastern Daylight Time US \(UTC-4\)](#)
- Place: Google-hangout, [https://plus.google.com/hangouts/\\_/event/c1glu6soq43r1rr6ou17qtobug8](https://plus.google.com/hangouts/_/event/c1glu6soq43r1rr6ou17qtobug8)

## Attendees

- Andrew Woods
- Esme Cowles
- Chris Beer
- Daniel Davis ★
- Declan Fleming
- A. Soroka
- Benjamin Armintor
- ~~Zhiwu Xie~~
- Neil Jefferies
- Yinlin Chen

## Agenda

1. Clarify objectives of this group (reference: [charter](#))
2. Establish approach for addressing initial objectives
  - a. Validate components and APIs
  - b. Identify any questions or concerns about Fedora 4.0 architecture in the context of the project's goals, objectives, and use cases
  - c. Assess core vs. modular features
3. Thought exercise: what would be the technical "risks" of releasing 4.0 Production \*now\*?

## Discussion

### Group objectives

1. Help figure out if Fedora 4 (as developed) is meeting user requirements
  - a. i.e. Ensure F4 platform is situated to address community needs going forward
2. Summarize group's findings in a report
  - Note: term 3-months, mid-Aug - mid-Oct
  - Note2: it would be nice to have some tech continuity in making decisions




## Approach

### Items for team assessment

1. REST API
  - Are immediate updates required?
  - We should version the API independently
    - a. This offers multiple backend implementations/optimizations
2. Performance
  - a. Read
  - b. Writes
    - i. Many small files
    - ii. Large files
    - iii. High throughput
  - c. Scalable serialization to disk
    - Need to measure scale of load that async serialization can meet
    - Need to clarify async approaches: messaging and sequencers
  - d. Replication of objects to another repository instance
  - e. Full re-indexing
  - f. Full integrity checks
3. Multi-node / Clustered configurations
  - a. High availability
  - b. Bulk ingest
  - c. High read loads
    - Note: generally need to define what clustering provides
4. ModeShape
  - a. Assess persistence approach (i.e. bit-level object and datastream persistence)
5. Preservation-worthiness

# Actions

## For next week

1.  All: bring expectations and goals (grounded in production use cases)
  - Not just performance but all qualities/functions
  - ... Goals that improve upon Fedora 3
  - ... And/Or goals that were supported by Fedora 3 and still need to be supported by Fedora 4
2.  Esmé: update [architecture diagram](#)
3.  Andrew: bring summary of performance tests/results to-date

## Ongoing

1. Compare F3 and F4 ingest rates
2. Revisit basic principles: "Why repository? Why Fedora?"