

DRAS-TIC FEDORA UPDATE

Using Data Partitions and Stateless Servers to
Scale Up Fedora Repositories

Gregory N. Jansen
University of Maryland iSchool
IEEE Big Data 2019: 4th Computational Archival Science (CAS) Workshop
Los Angeles, 11th December 2019

Overview

- Need: Scale up Storage and Performance at Web Scale
- Tech: Stateless Servers and Distributed Database
- Approach: Partner, Build, and Test
- Results: Graphs and comparisons
- Fine Tuning
 - Specification Compliance Work
 - Docker Stack Development
- Future Work
 - UMD Libraries
 - National Parks Service Capital Region
 - Distributed Compute Demonstrations



Community Partners - Year 3

Smithsonian Institution


Amherst College -- Inrupt.com

University of Illinois
at Urbana-Champaign

Georgetown University

University of Maryland Libraries

National Parks Service*



Storage and Performance Needs



Univ. of Maryland Libraries

- Predictable costs
- No storage limits
- Maintain performance
- Fewer repositories
- New A/V, science collections



Smithsonian

The Smithsonian Institution

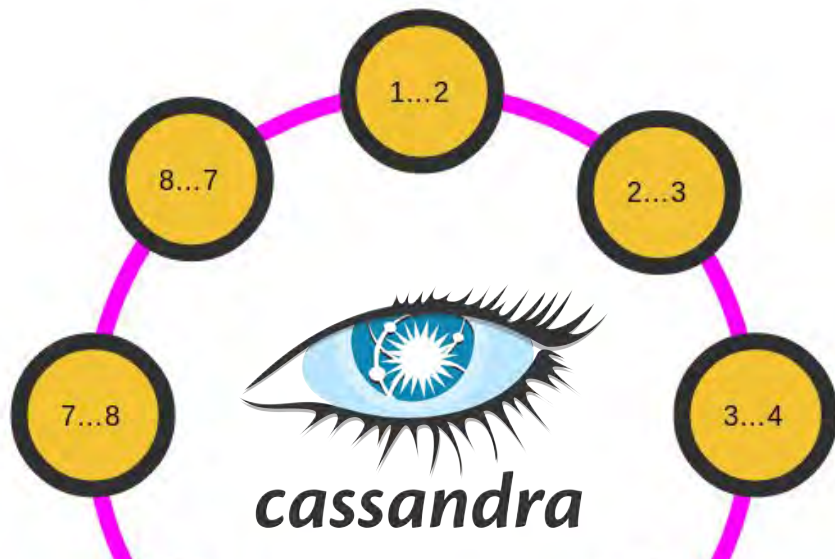
- Large and diverse collections of scientific and cultural data
- Existing science workflow tools, i.e. genomics desktop software
- Compute on repository objects

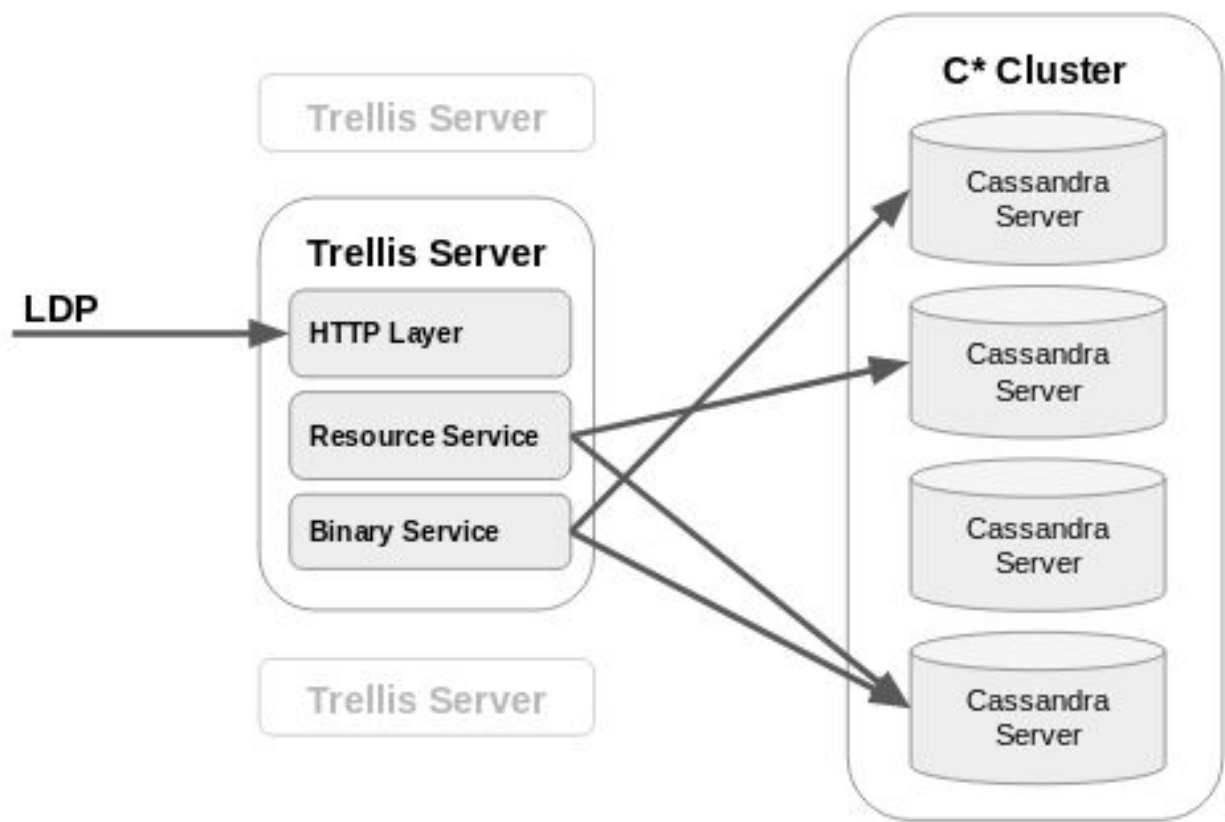
Stateless Servers

- Zero coordination btw. web servers
- Add more servers to meet more load
- Use cloud elasticity

Distributed Database

- Fast writes with tunable consistency
- No transaction locking
- Replication and repair is built in
- Add more servers to increase storage (Apple ran 75,000)





Approach: Partner, Build, and Test

Trellis Linked Data Platform

Developed by Aaron Coburn while at Amherst College (now Inrupt.com)

Stateless server implementing LDP, Memento, WebAC

Platform adopted by Inrupt.com - TBL and Social Linked Data (Solid)

Version 0.9 to be released soon

Cassandra Storage Module

Largely contributed by Adam Soroka at the Smithsonian Institution

Now maintained as a part of Trellis LDP Extensions Project

Performance Testbed at UMD iSchool Digital Curation Innovation Center

Fedora API Testing Partnership with Fedora Community and Lyrasis

DRAS-TIC Testbed Website

DRAS-TIC Testbed at UMD

Digital Repositories at Scale that Invite Computation

Purpose

Our goal is to prove the performance of digital repository architectures under emerging next-generation workloads. We do this by the building and testing various implementations of the Linked Data Platform (LDP), Memento, and Fedora 5 family of repository specifications. Testing workloads are based on systems requirements gathered in early 2018 from our institutional partners.

Studied Systems

- Trellis LDP (Cassandra)
 - ◊ [Website](#), [Repository](#)
 - ◊ [Fedora API Test Suite: Report](#), [TestNG Detail](#)
- Trellis LDP (PostgreSQL)
 - ◊ [Website](#), [Repository](#)
 - ◊ [Fedora API Test Suite: Report](#), [TestNG Detail](#)
- Fedora 5 Reference Implementation (TBA)
 - ◊ [Website](#), [Repository](#)

Institutional Partners

- [Summary of Partner Requirements](#)
- [University of Maryland Libraries - interview notes](#)
- [Georgetown University Libraries - interview notes](#)
- [The Smithsonian Institution \(co-creators of Trellis Cassandra\) - interview notes](#)
- [Amherst College \(creators of Trellis LDP\)](#)
- [National Center for Supercomputing Applications \(creators of Brown Dog\)](#)

Publications

- [DRAS-TIC Linked Data: Evenly Distributing the Past](#)

Use Cases






Digital Manuscripts Collection
Big Scientific Dataset
Citizen Science Image Curation
Social Media Archives

Performance Workloads

- Large file ingest
- Deeply nested collection ingest and update
- Very large folders ingest and update (AKA "Supernode")
- Incremental scaling of storage nodes
- Dynamic scaling of front-end nodes
- Pervasive metadata and feature extraction (Brown Dog)
- On-demand conversion to access formats (Brown Dog)

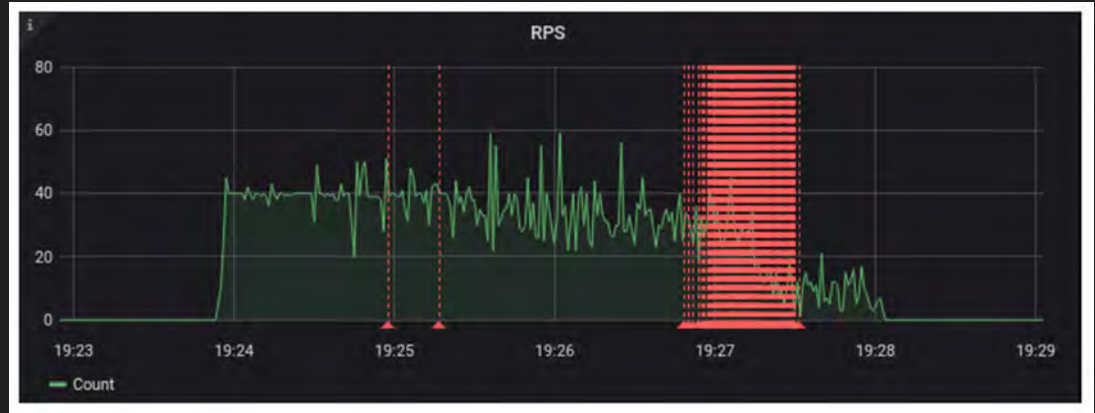
Table of Tests with Links to Dashboards

Test Results


Dashboard	Date	LDP Docker Name	LDP Docker Tag	LDP Code Commit	Test Script	Test Code Commit	Running Nodes	Flipped Nodes	Back-end Nodes	C* Replication	C* Max Chunk Size	C* Binary Read Consistency	C* Binary Write Consistency	C* RF Read Consistency	C* RF Write Consistency	MS Events Enabled	Test Code Repo	LDP Code Repo	OK	KO	AvG	Min
 comment	2019-11-20	trellisdp/trellis-cassandra	develop		ldp.StressTestIngest	dd2f566	1	4	4							false	git@github.com:UMD-DCI/galling-testbed.git		3983	0	67	4070
 comment	2019-11-20	trellisdp/trellis-cassandra	develop		ldp.StressTestIngest	dd2f566	4	4	4							false	git@github.com:UMD-DCI/galling-testbed.git		8886	3933	19703	60065
 comment	2019-10-08	trellisdp/trellis-cassandra	test	trellis-cassandra-0.8-110-ge0c5436	ldp.StressTestIngest	dd2f566	4	4	4	1	1048576	ONE	ONE	ONE	ONE	false	git@github.com:UMD-DCI/galling-testbed.git	git@github.com:trellis-ldp/trellis-cassandra.git	9954	3004	907	5292
 comment	2019-10-03	trellisdp/trellis-cassandra	test	trellis-cassandra-0.8-84-g0a5ffc3	ldp.StressTestIngest	dd2f566	4	4	4	1	1048576	ONE	ONE	ONE	ONE	false	git@github.com:UMD-DCI/galling-testbed.git	git@github.com:trellis-ldp/trellis-cassandra.git	15950	0	40	4248
 comment	2019-10-02	trellisdp/trellis-cassandra	test	trellis-cassandra-0.8-109-g1204735-dirty	ldp.StressTestIngest	dd2f566	4	4	4	1	1048576	ONE	ONE	ONE	ONE	false	git@github.com:UMD-DCI/galling-testbed.git	git@github.com:trellis-ldp/trellis-cassandra.git	9695	3209	1567	16151

Performance Testbed Results


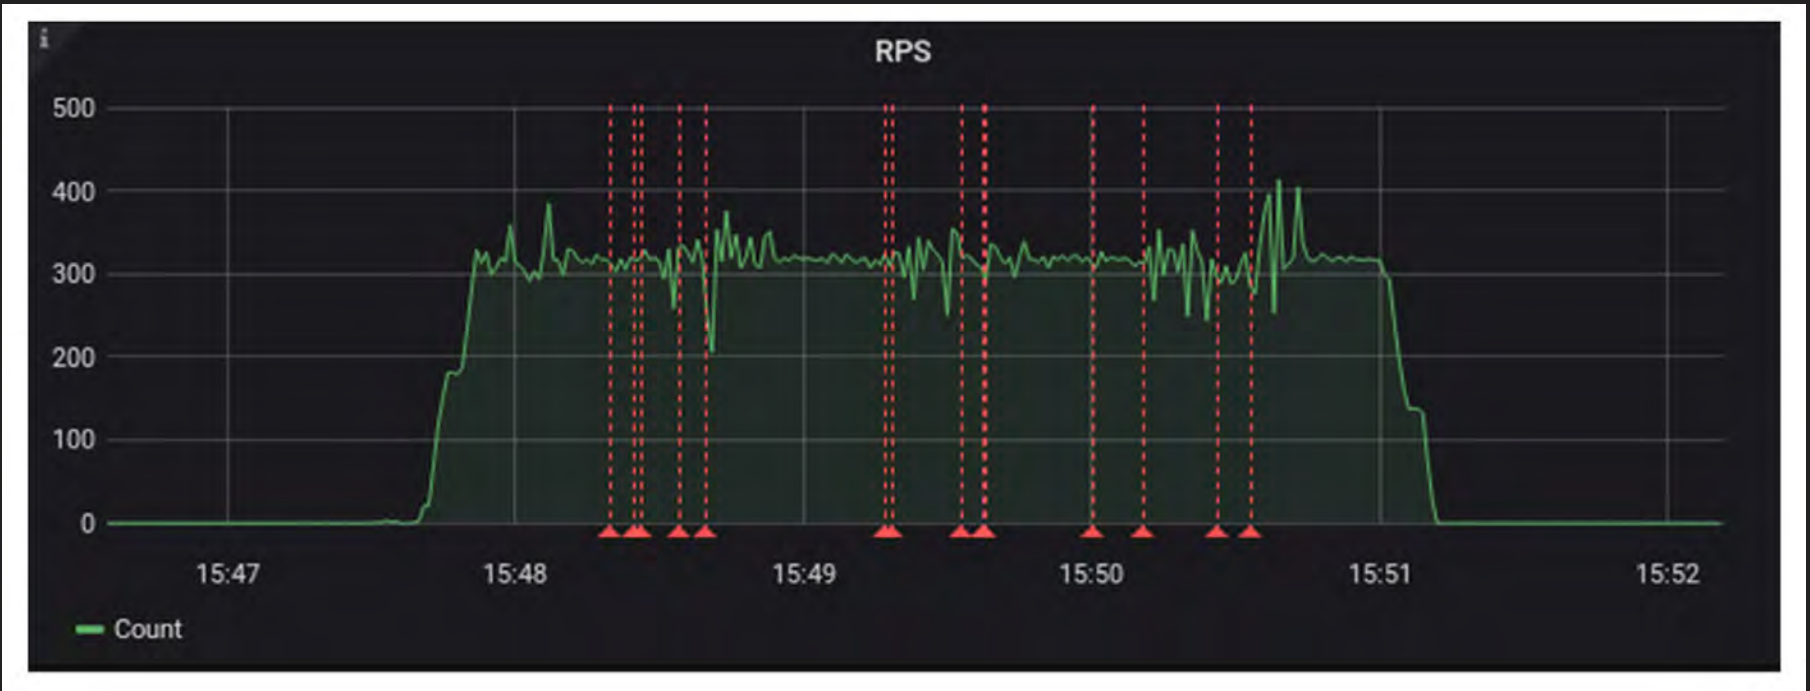
Fedora PostgreSQL degrades at 40 requests per second.



A 4 Trellis, 4 Cassandra stack handles 480 requests per second without failures.



Performance Testbed Results



Fedora API Test Suite Results

67 % of Fedora API Tests are now passing.

Most remaining test failures represent implementation decisions by Aaron Coburn, due to security or specification concerns.

DRAS-TIC Fedora is..

- **addressing API gaps in a Trellis extension module**
- **suggesting minor revisions to Fedora API Specification**
- **contributing features and bug fixes to the Fedora API Test Suite**

- **committed to making Trellis LDP a viable alternative for Fedora users**

Work Ahead

Develop and test a Docker-based repository stack composed of:

- Trellis LDP
- Apache Cassandra
- DRAS-TIC Fedora Extension Module
- Quarkus Java MicroProfile Application Container

Test UMD Libraries' workflow tools against this stack. (Ingest works.)

Basis for a National Parks Service repository prototype.

Thank You

Questions?

<https://drastic-testbed.umd.edu>

<https://trellisldp.org>