THE DIGITAL PRESERVATION NETWORK

Architecture, Process and Policies

Open Repositories July 10, 2013

David Minor – Chronopolis Program Head
Scott Turnbull – Lead Engineer for the Academic Preservation Trust

The Digital Preservation Network (DPN) was formed to ensure that the complete scholarly record is preserved for future generations.

What Is DPN?

Technical staff and systems from extant large- scale preservation repositories … working with groups of experts in: succession rights, business services, communications and research data.

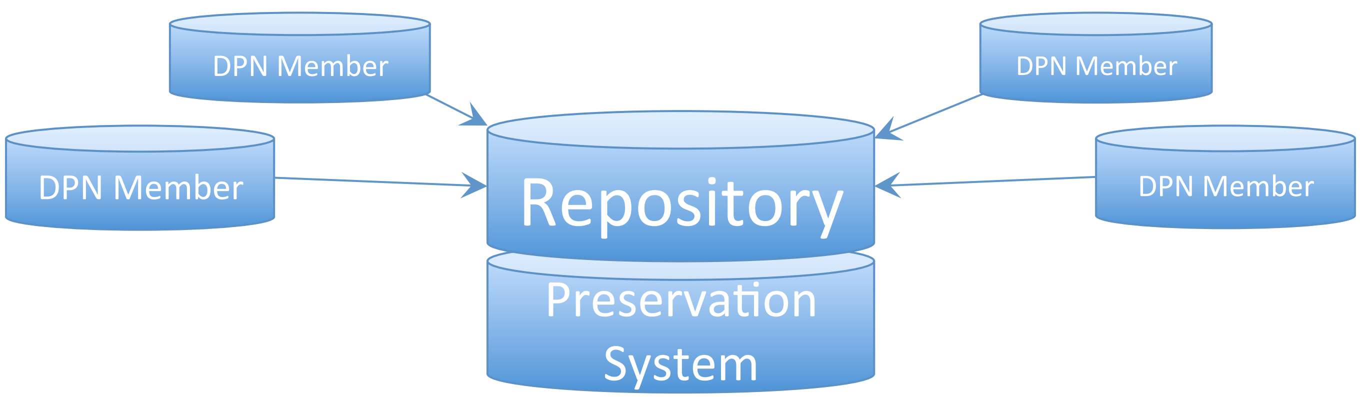

DPN Benefits

1. Resilience
2. Succession
3. Economies of scale
4. Efficiency
5. Extensibility
6. Security

What Does DPN Do?

1. Establishes a network of heterogeneous, interoperable, trustworthy, preservation repositories (Nodes)
2. Replicates content across the network, to multiple nodes
3. Enables restoration of preserved content to any node in the event of data loss, corruption or disaster
4. Ensures the ongoing preservation of digital information from depositors in the event of dissolution or divestment of depositors or an individual repository
5. Provides the option to (technically and legally) "brighten content" preserved in the network over time

Critical Assumptions

- All content enters DPN by deposit into one of the DPN Nodes, known as a “First Node.”
- Nodes with copies of this content are “Replicating Nodes”.
- DPN Members will work directly with First Nodes to negotiate contracts and determine service levels
- Service levels and contracts will reflect “standard” DPN services; they may also reflect the First Node’s unique offerings in terms of access, hosting or other services.
- Content in Replicating Nodes will be held “dark”, and inaccessible except for preservation actions.

Critical Assumptions

- DPN shall redistribute preserved content as Nodes enter and leave the Network, ensuring continuity of preservation services over time.
- DPN will provide a large-scale network of dark archives that enable the opportunity to brighten content in the future, but does not mandate how this is done.
- Depositors, First Nodes and their designated communities will collaborate to ensure that the information contents of DPN deposits are accessible for reuse in the future, using the appropriate (and evolving) community standards for any given set of content.

Initial technical partners

Initial DPN launch will feature five nodes:

- Academic Preservation Trust (APTrust)
- Chronopolis
- HathiTrust
- Stanford Digital Repository (SDR)
- University of Texas Data Repository (UTDR)

And a participating partner:

- DuraSpace

DPN Usage Scenarios

Scenario 1: Ingest & Replication

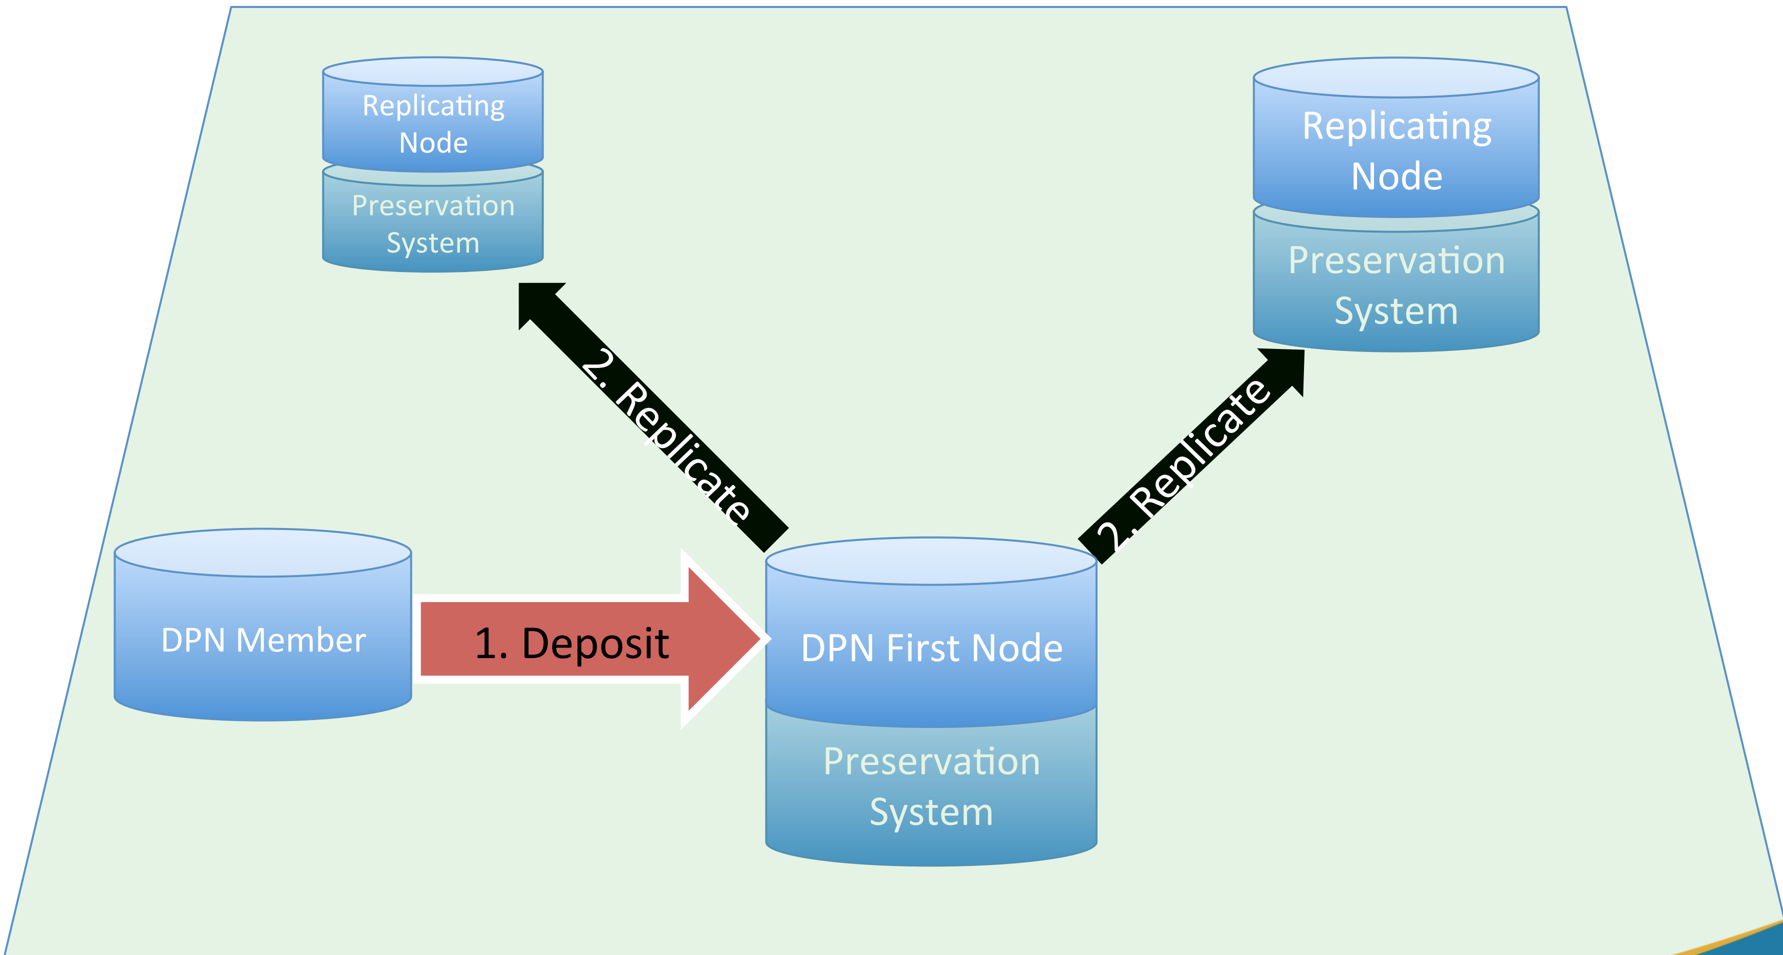

Scenario 2: Restoration of Content

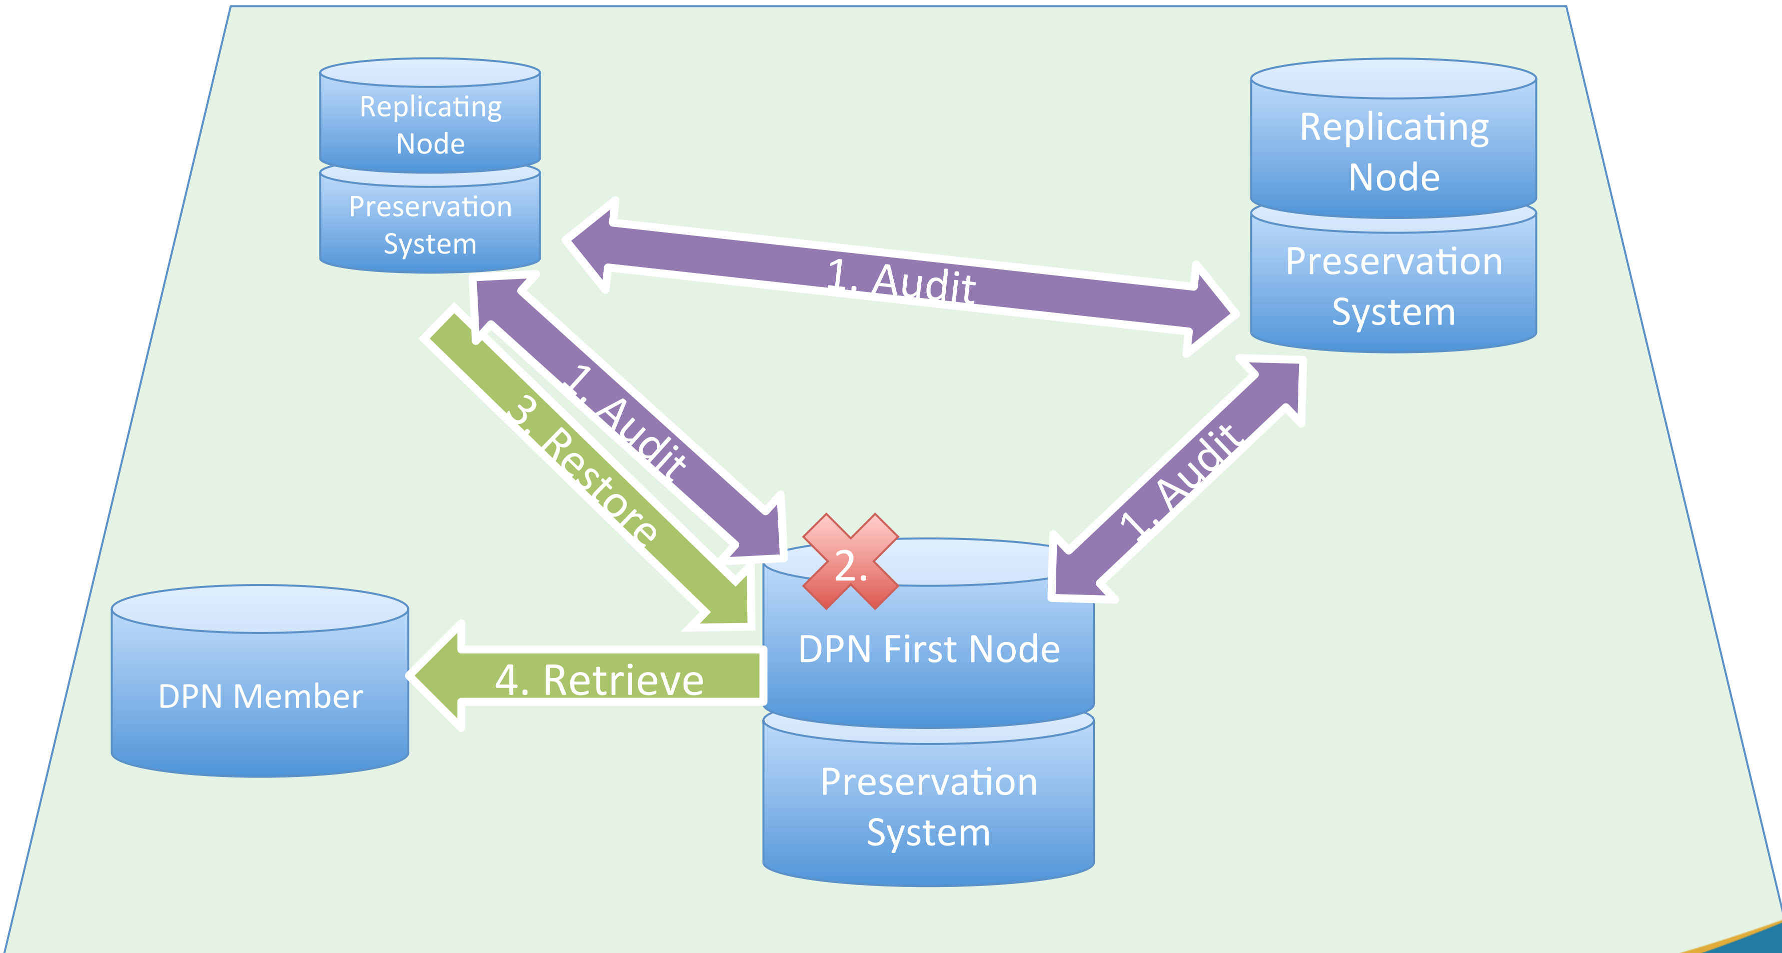

Scenario 3: First Node Cessation

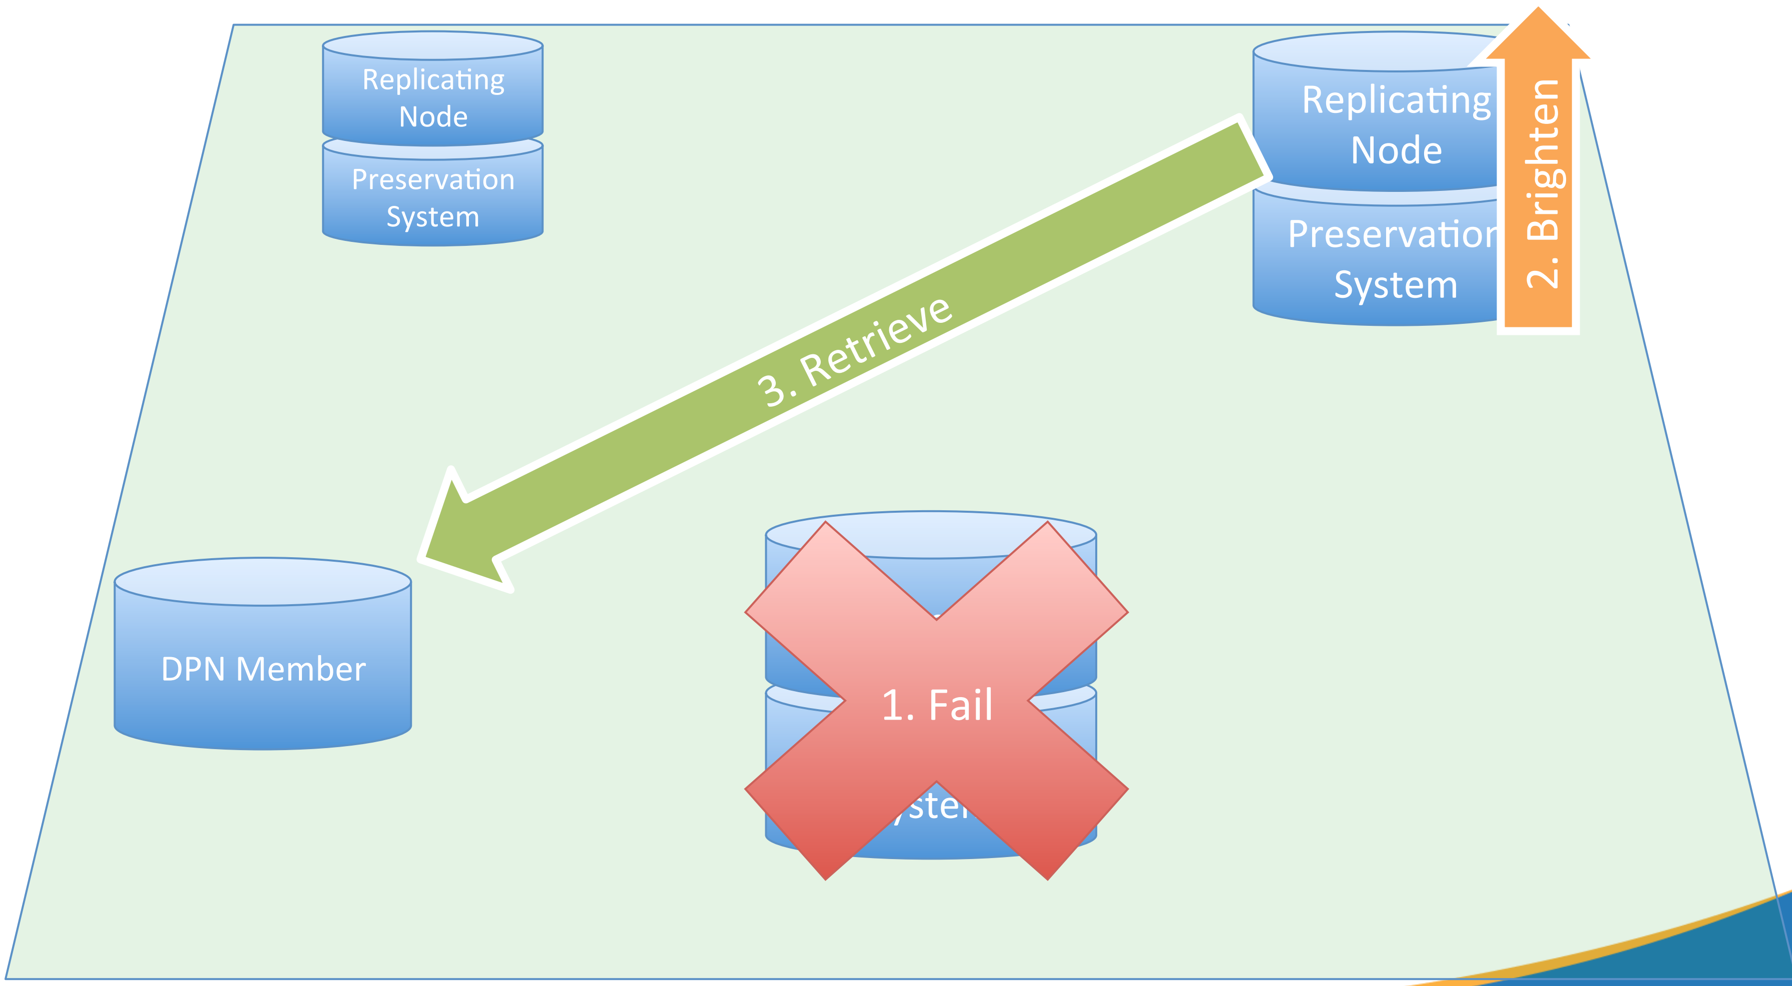

Scenario 4: Successioning

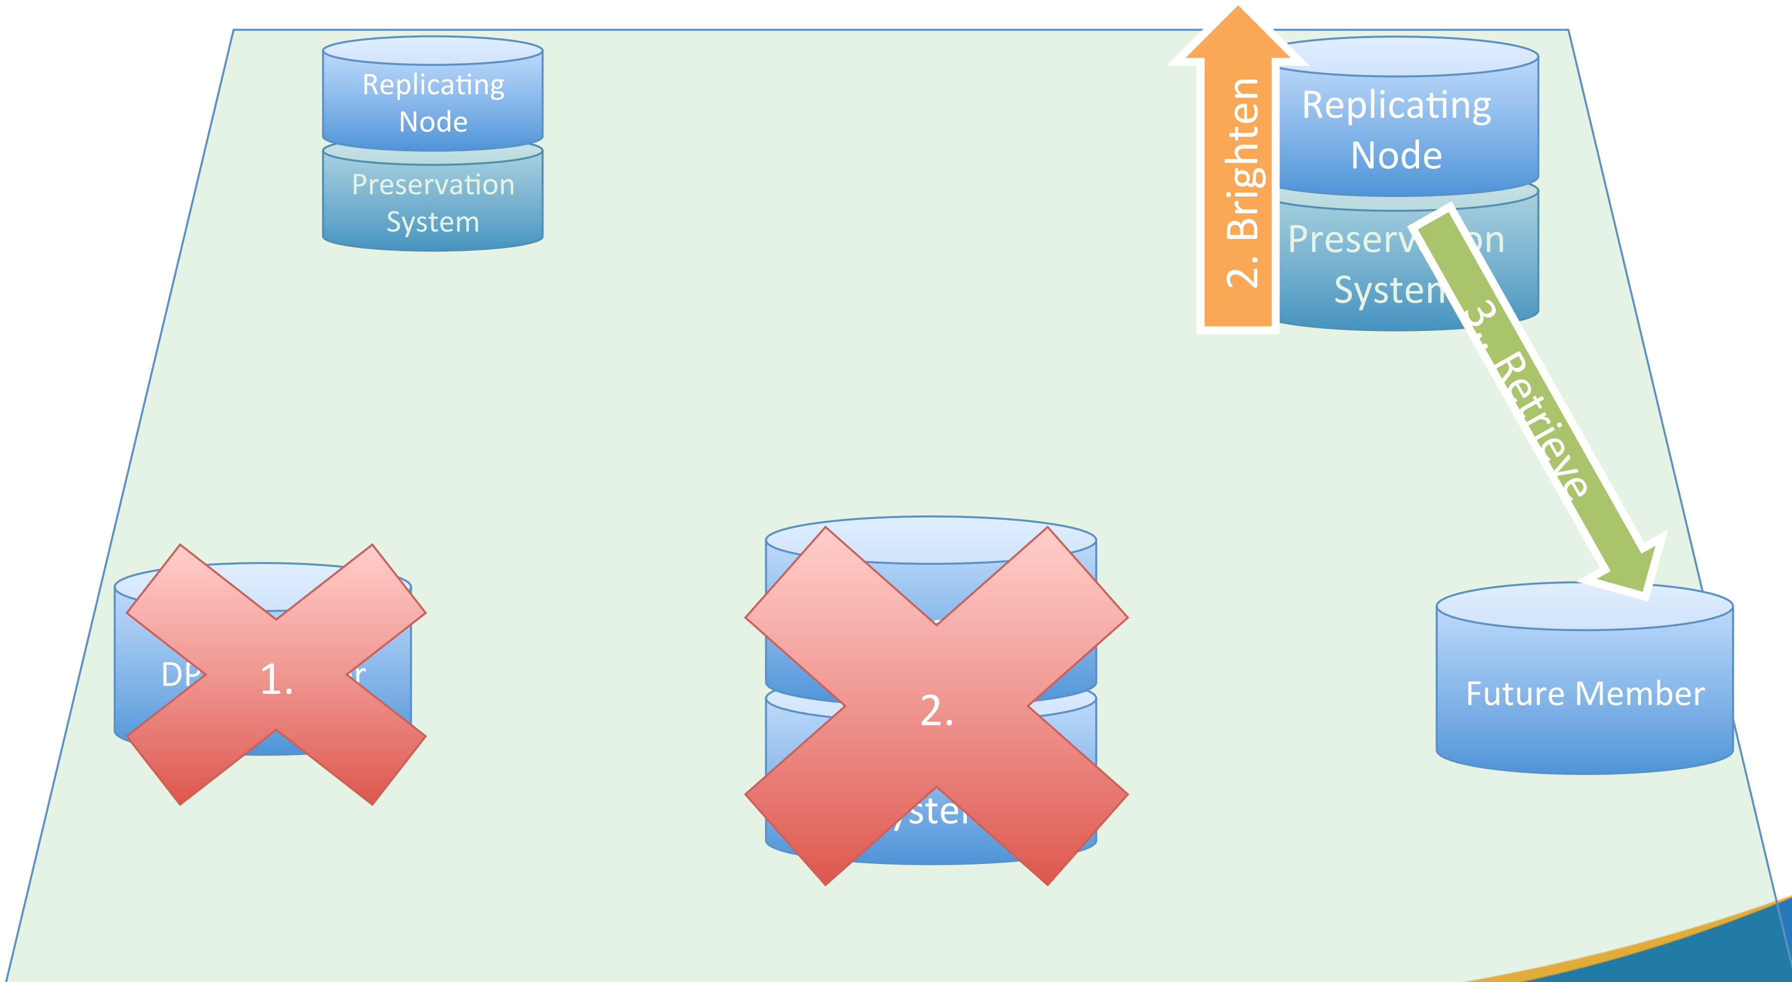

Architectural Overview

Architectural Overview

- Architectural Premise
  - *Core capabilities founded on proven institutions and repositories*
- Design Considerations
  - *Distributed Nodes, loosely coupled*
  - *Standards and protocol-based integrations*
  - *Separate implementations*
  - *Distributed infrastructure*

Top-level Architecture

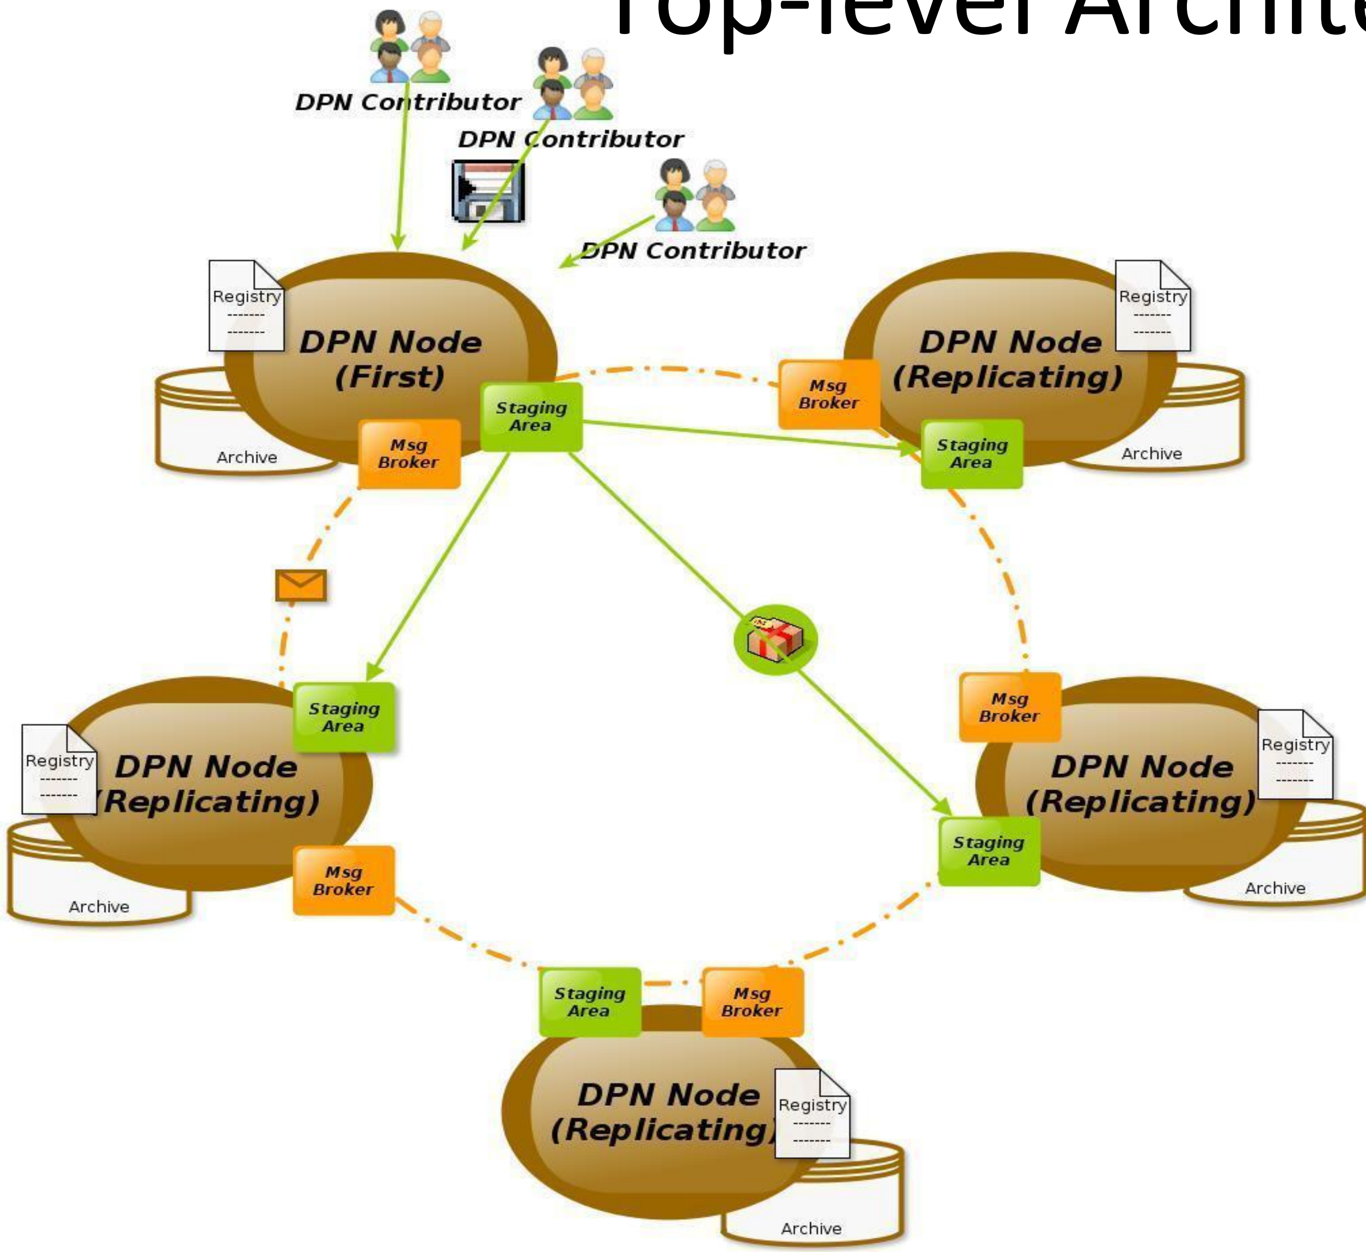

Infrastructure Components

- *Archive/Repository*
- *Federated Messaging*
- *Distributed Registry*
- *Transfer Mechanisms*
- *Content Packaging*
- *Security and Encryption*

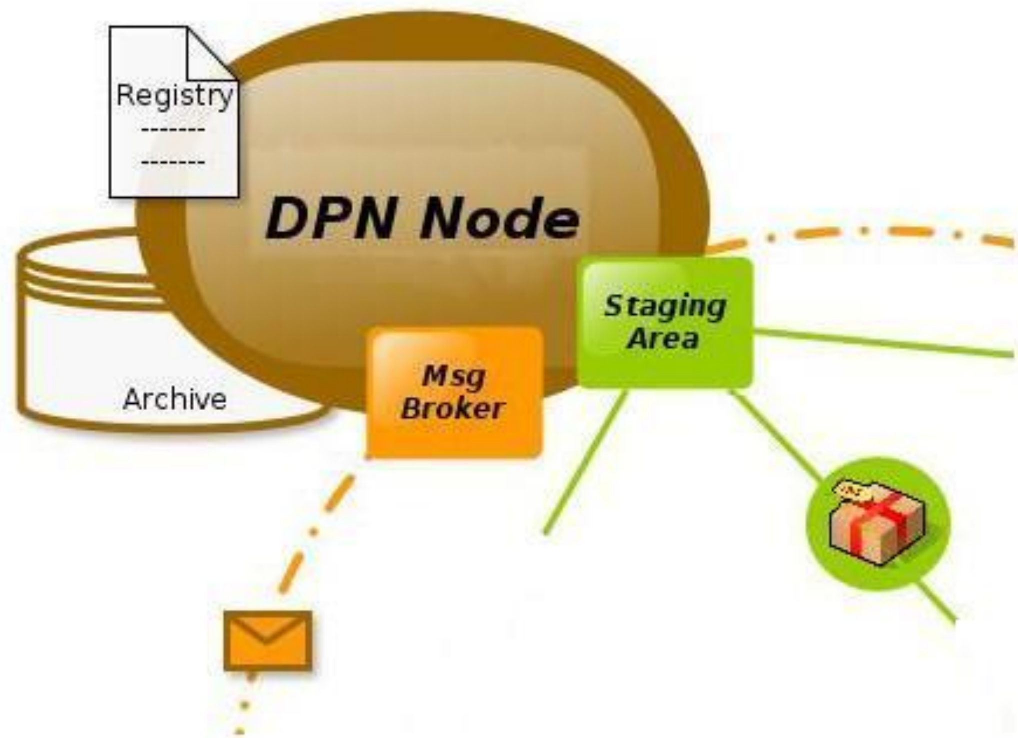

Infrastructure Components

- ***Archive/Repository***
- *Federated Messaging*
- *Distributed Registry*
- *Transfer Mechanisms*
- *Content Packaging*
- *Security and Encryption*

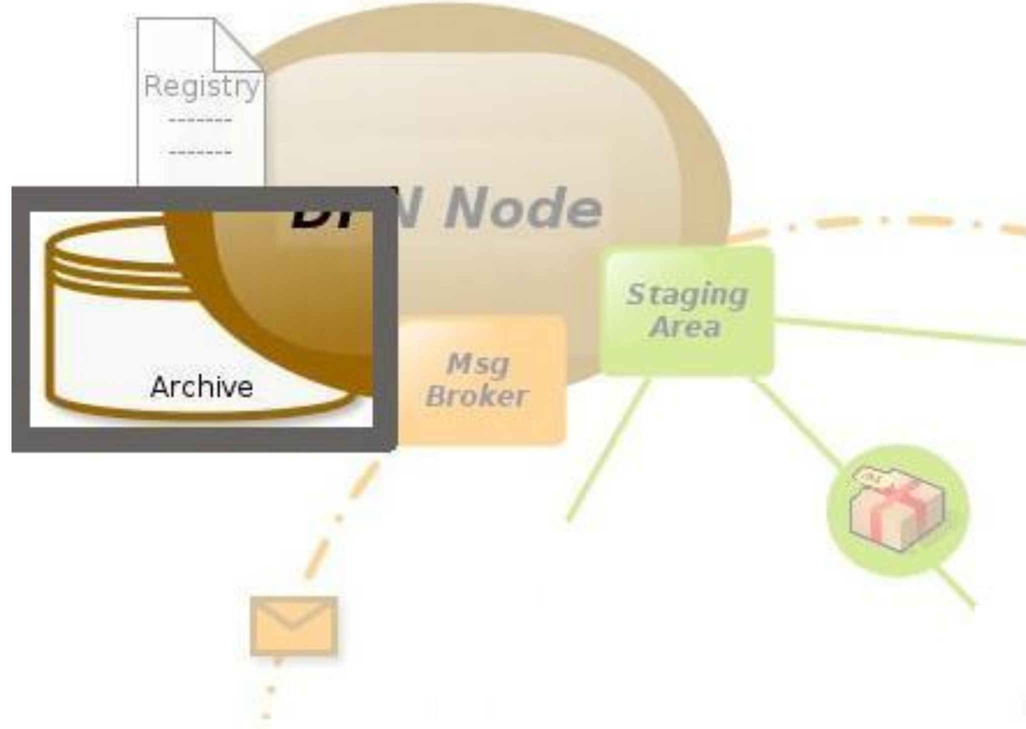

Infrastructure Components

- *Institutional Archive/Repository*
- ***Federated Messaging***
- *Distributed Registry*
- *Transfer Mechanisms*
- *Content Packaging*
- *Security and Encryption*

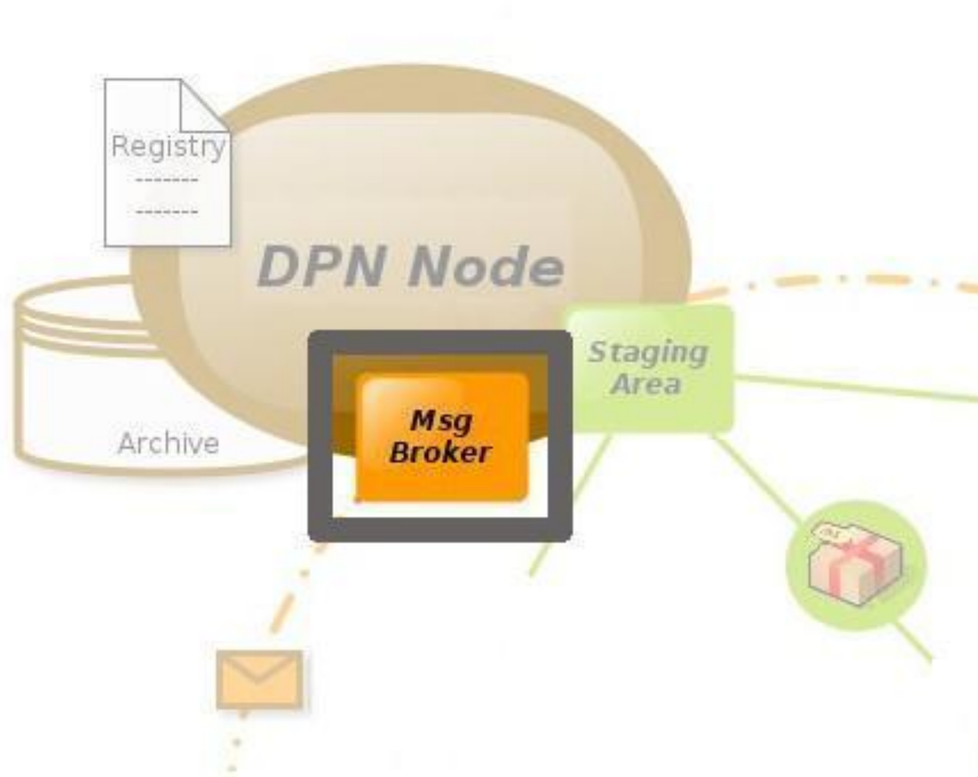

Federated Messaging

- DPN uses messaging for in-band communication
- Using RabbitMQ message brokers, which support AMQP (Advanced Messaging Queueing Protocol)
- RabbitMQ also supports federated messaging easily via a plugin
- DPN messaging model uses Topic Queues for broadcast messages and direct queues for one-to-one communication between nodes

Messaging Model

- Broadcast messages are sent to all node brokers
- Node brokers federate all messages, so if one broker is down it is still possible to communicate

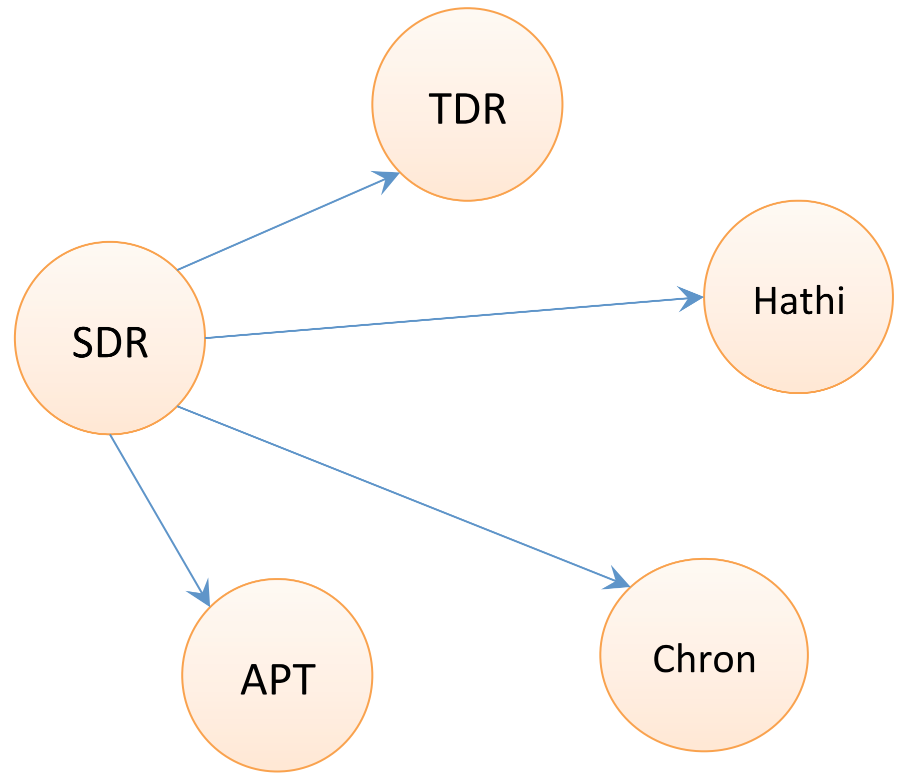

Messaging Model

- Direct messages are between two nodes, used for replies in a message sequence
- Broker federation still applies, so communication channels are redundant

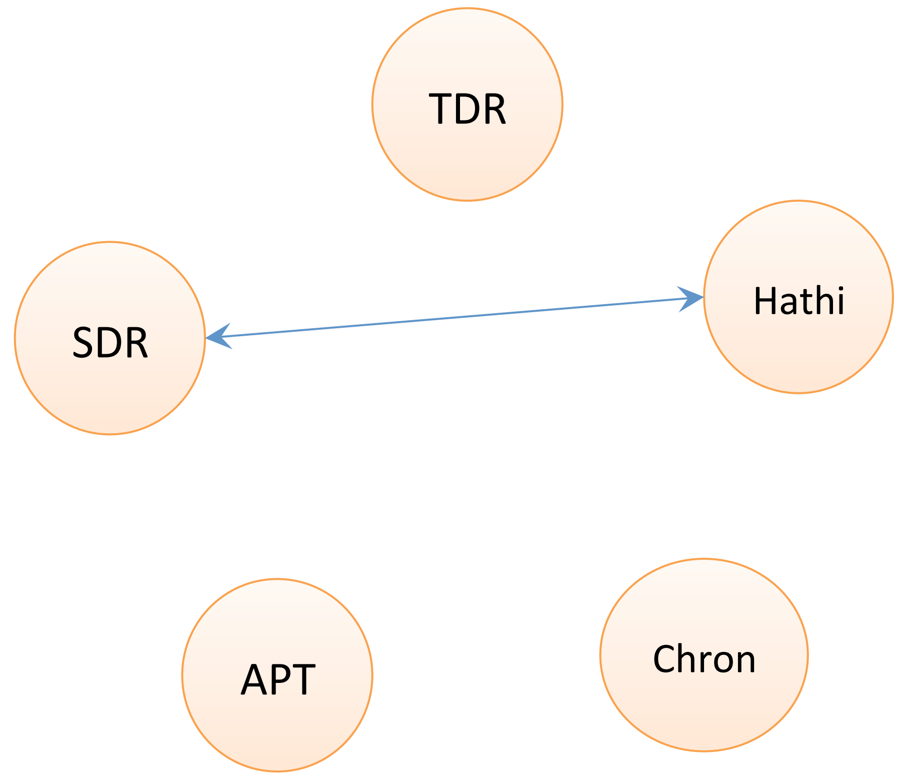

Messaging Control Flows

- Message control flows are transactional, and asynchronous

Examples of control flows

- Replication Request
- Registry Item Create
- Registry Synchronize
- Recovery (digital object, registry entry, registry, etc.)
- Fixity Audit flows

At any given time each node may be handling multiple message control flows/ sequences at once

Infrastructure Components

- *Institutional Archive/Repository*
- *Federated Messaging*
- ***Distributed Registry***
- *Transfer Mechanisms*
- *Content Packaging*
- *Security and Encryption*

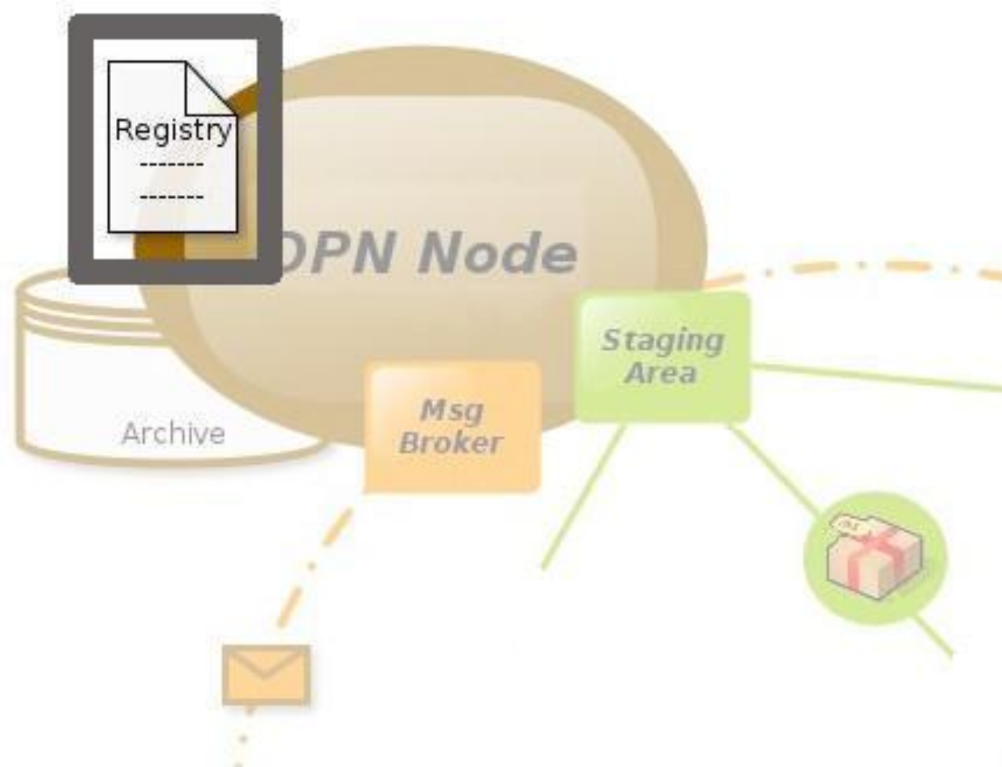

Registry

- Messages support Registry services
  - Create, read, update, delete
  - Delete is a special case, with special handling
- Creation of a new Registry entry will be at the request of a First Node.
  - It will only happen after a quorum of correct copies have been made to Replicating Nodes
  - The Registry entry will be updated at ALL nodes
  - Note that this is a distributed environment, so we expect that the registries will be eventually consistent following Brewers theorem

Federated Registry Synchronization

- At any given time there is a possibility that a node is down, and may not receive Registry messages to create entries, or update entries
- The First Node that issues a create/update can wait and retry, but eventually may give up, i.e. the time to live for the message expires
- To accommodate synchronization, each node will keep a list of registry entries that the node has updated within its own registry
- During a synchronization, each node will exchange synchronization lists and compare against its own list, items missing will show up on other nodes lists and can be recovered

Infrastructure Components

- *Institutional Archive/Repository*
- *Federated Messaging*
- *Distributed Registry*
- ***Transfer Mechanisms***
- *Content Packaging*
- *Security and Encryption*

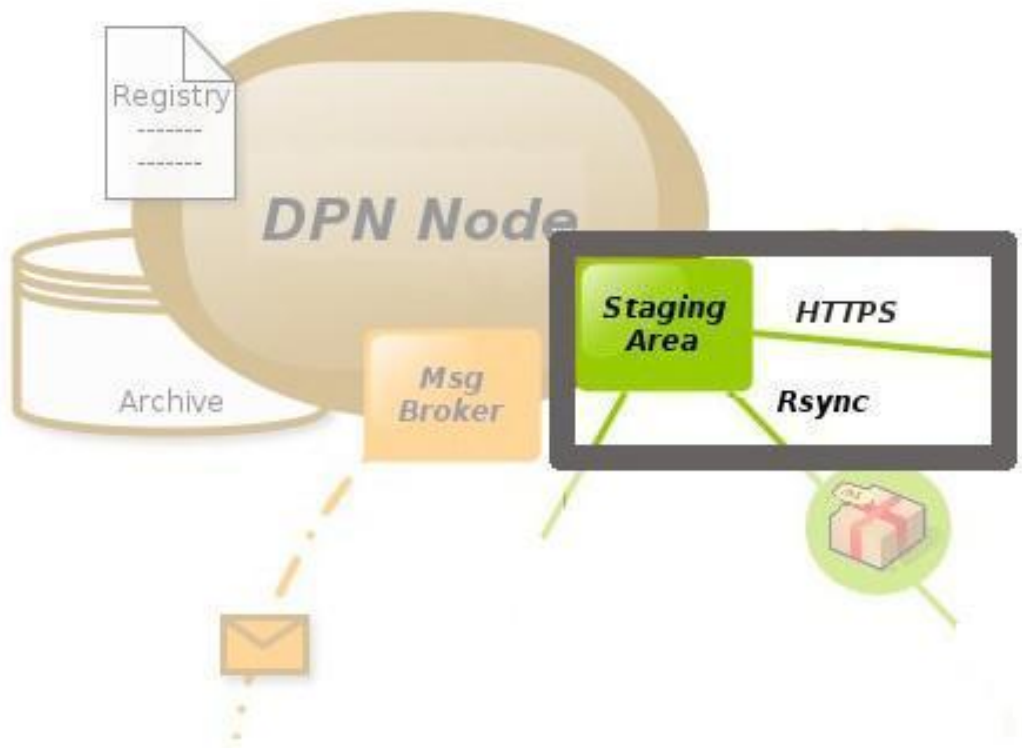

DPN Data Transport

- Used only for copying bags within DPN:
  - Initial Replication
  - Restoring replicas after a failure
- Use widely-supported, easy-to-script transport mechanisms
- Also plan to support high-performance mechanisms where possible
- Confirmation of fixity done outside transport

DPN Transport Mechanisms

- Work in progress - not a final list
- HTTPS
  - Simple to use, widely supported
- rsync-over-ssh
  - Ubiquitous on Unix hosts
- GridFTP - More technical complexity, but excellent for long, fat pipes

Infrastructure Components

- *Institutional Archive/Repository*
- *Federated Messaging*
- *Distributed Registry*
- *Transfer Mechanisms*
- ***Content Packaging***
- *Security and Encryption*

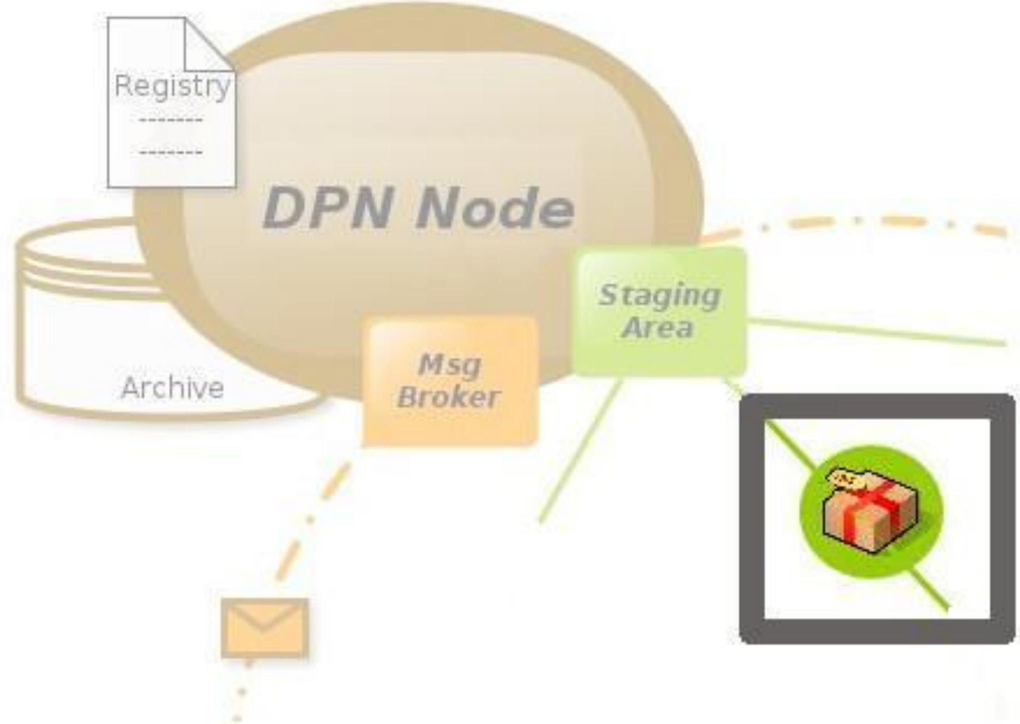

DPN Packaging - BagIt

- Standard packaging method shared by all nodes
- Minimal, standard bag metadata to enable tracking identity, source and fixity of content bags
- No DPN-wide requirements on descriptive metadata or content structure below the top level bag

DPN Packaging - BagIt

- DPN packages will conform to the BagIt packaging format
- DPN packages may either be
  - serialized (e.g. a single tar)
  - un-serialized (e.g. exploded directory structure)
- DPN packages will conform to a TBD BagIt profile, still under discussion

Infrastructure Components

- *Institutional Archive/Repository*
- *Federated Messaging*
- *Distributed Registry*
- *Transfer Mechanisms*
- *Content Packaging*
- ***Security and Encryption***

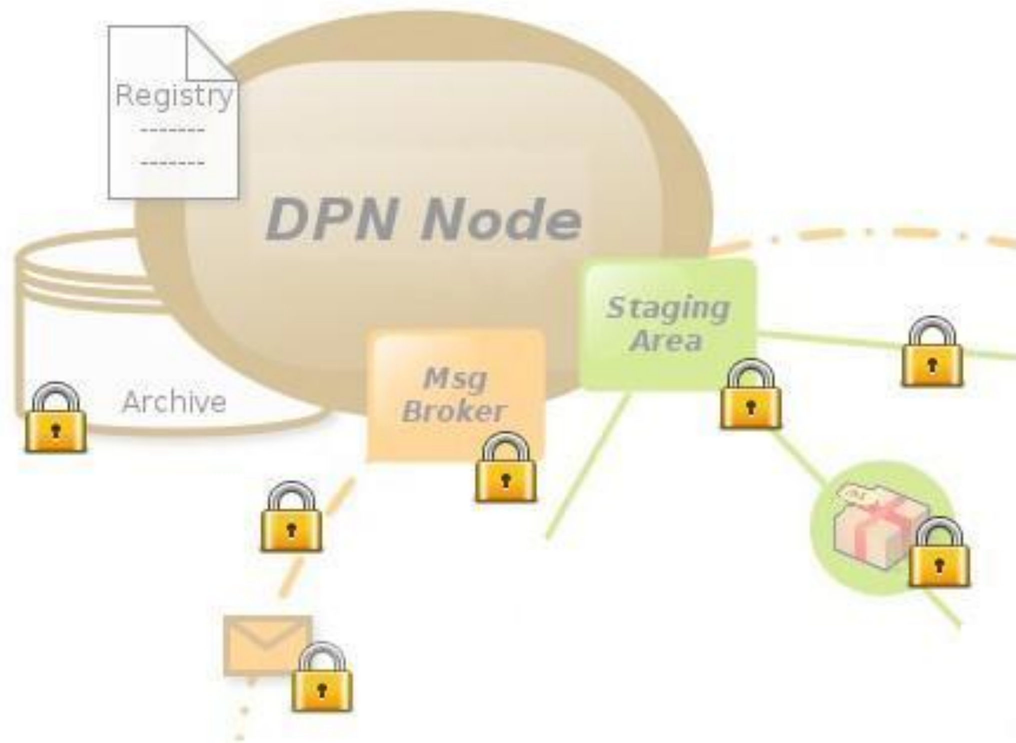

DPN Encryption *work in progress*

- Some content may be encrypted at rest
- Depositors / First Nodes must have confidence that content is secure
- Key escrow to allow content to survive any succession events

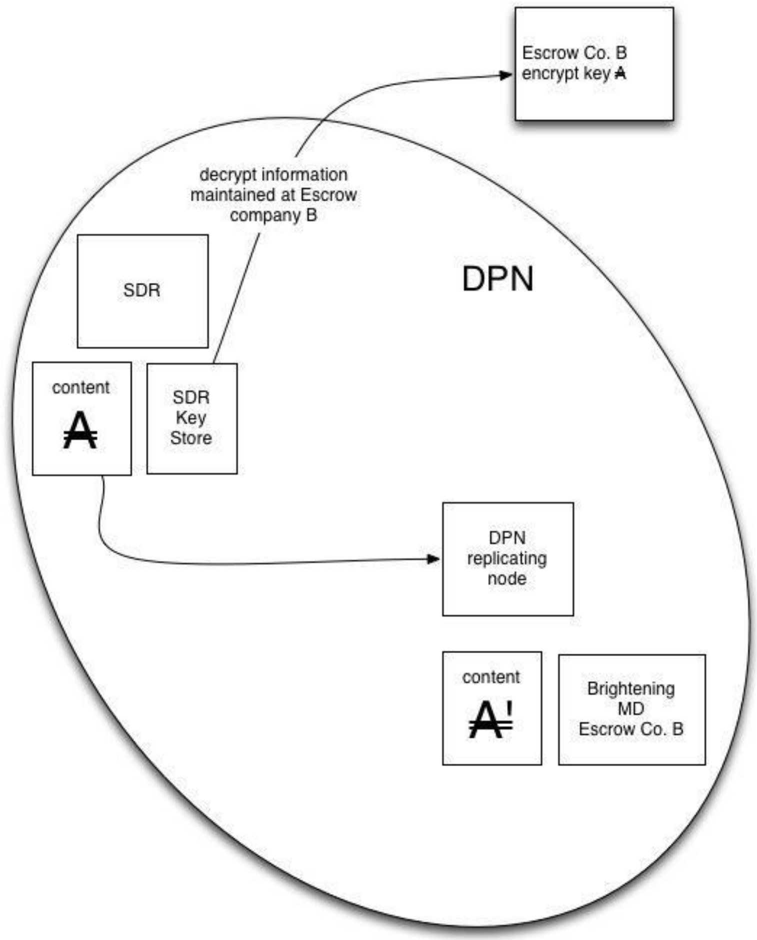

Development Paradigm

- Federated environment rather than a single application.
- Heterogeneity as a design principle in DPN means a different implementation at each Node.
- Open Standards vital for interaction between Federated Nodes.
- Heavy dependency on policy agreements shapes the conversation on standards.

Implementation Diversity

- APTrust - Python
- Chronopolis - Java
- HathiTrust - JRuby
- Stanford Digital Repository - Ruby
- University of Texas Data Repository - PHP
- Transfer protocols may vary per Node:
    - HTTPS
    - Rsync
    - Others perhaps

Concurrent Development

- Strong specifications are critical given diversity of implementations.
- GitHub for more social coding, code review tools, and tracking of changes over time.
- Consensus-based decision making by implementation team.
- Healthy debate over details of specifications have had very good results.

Development Challenges

- Diversity of architecture complicates growth of services and refactoring by # nodes.
- Diversity of missions between nodes in the federations make some implementation decisions more difficult to reach.
- Challenge coordinating a geographically diverse team with varying responsibilities and availability.

Advantages

- Diversity of nodes in the federation means we are able to draw on a pool of highly talented people.
- Right people in the right place with the right skills.
- Flexibility of Federation also avoids implementation conflicts that might otherwise occur.

DPN Charter Members

Arizona State
Brown
California Institute of Technology
Columbia
Cornell
Dartmouth
Duke
Emory
Harvard
Indiana
Iowa State
Johns Hopkins
Kansas State

MIT
Michigan State
NYU
Northwestern
NC State
Ohio State
Penn State Univ.
Purdue
Rutgers
Stanford
Syracuse
Texas A&M
Texas Tech Univ.
Tufts

Tulane
Univ. of Alabama
Univ. of Arizona
UC San Diego
Univ. of Chicago
Univ. of Florida
Illinois at Chicago
Illinois at Urbana-Champaign
Univ. of Iowa
Univ. of Kansas
Univ. of Kentucky
Univ. of Maryland
Univ. of Miami

Univ. of Michigan
Univ. of Minnesota
Univ. of Nebraska
Univ. of North Carolina
Univ. of North Texas
Notre Dame
Univ. of Tennessee
Univ. of Texas
Univ. of Virginia
Univ. of Washington
Univ. of Wisconsin
Vanderbilt Univ.
Virginia Tech
Yale Univ.

Texas Digital Library
California Digital Library
John D. Evans Foundation
American Council on Education

For more information ...

DPN Website:

http://www.dpn.org

DPN Public Wiki:

https://wiki.duraspace.org/display/DPNC/Digital+Preservation+Network

Contacts:

General: inquiry@dpn.org

Steven Morales, Project Manager: steven.morales@dpn.org

DPN

The Digital Preservation Network