METADATA AND FILESYSTEMS

(a brief synopsis of some of the presentation from Fall 2018 HPCUF meeting)

Paul Muzio

Chair, HPC User Forum Steering Committee

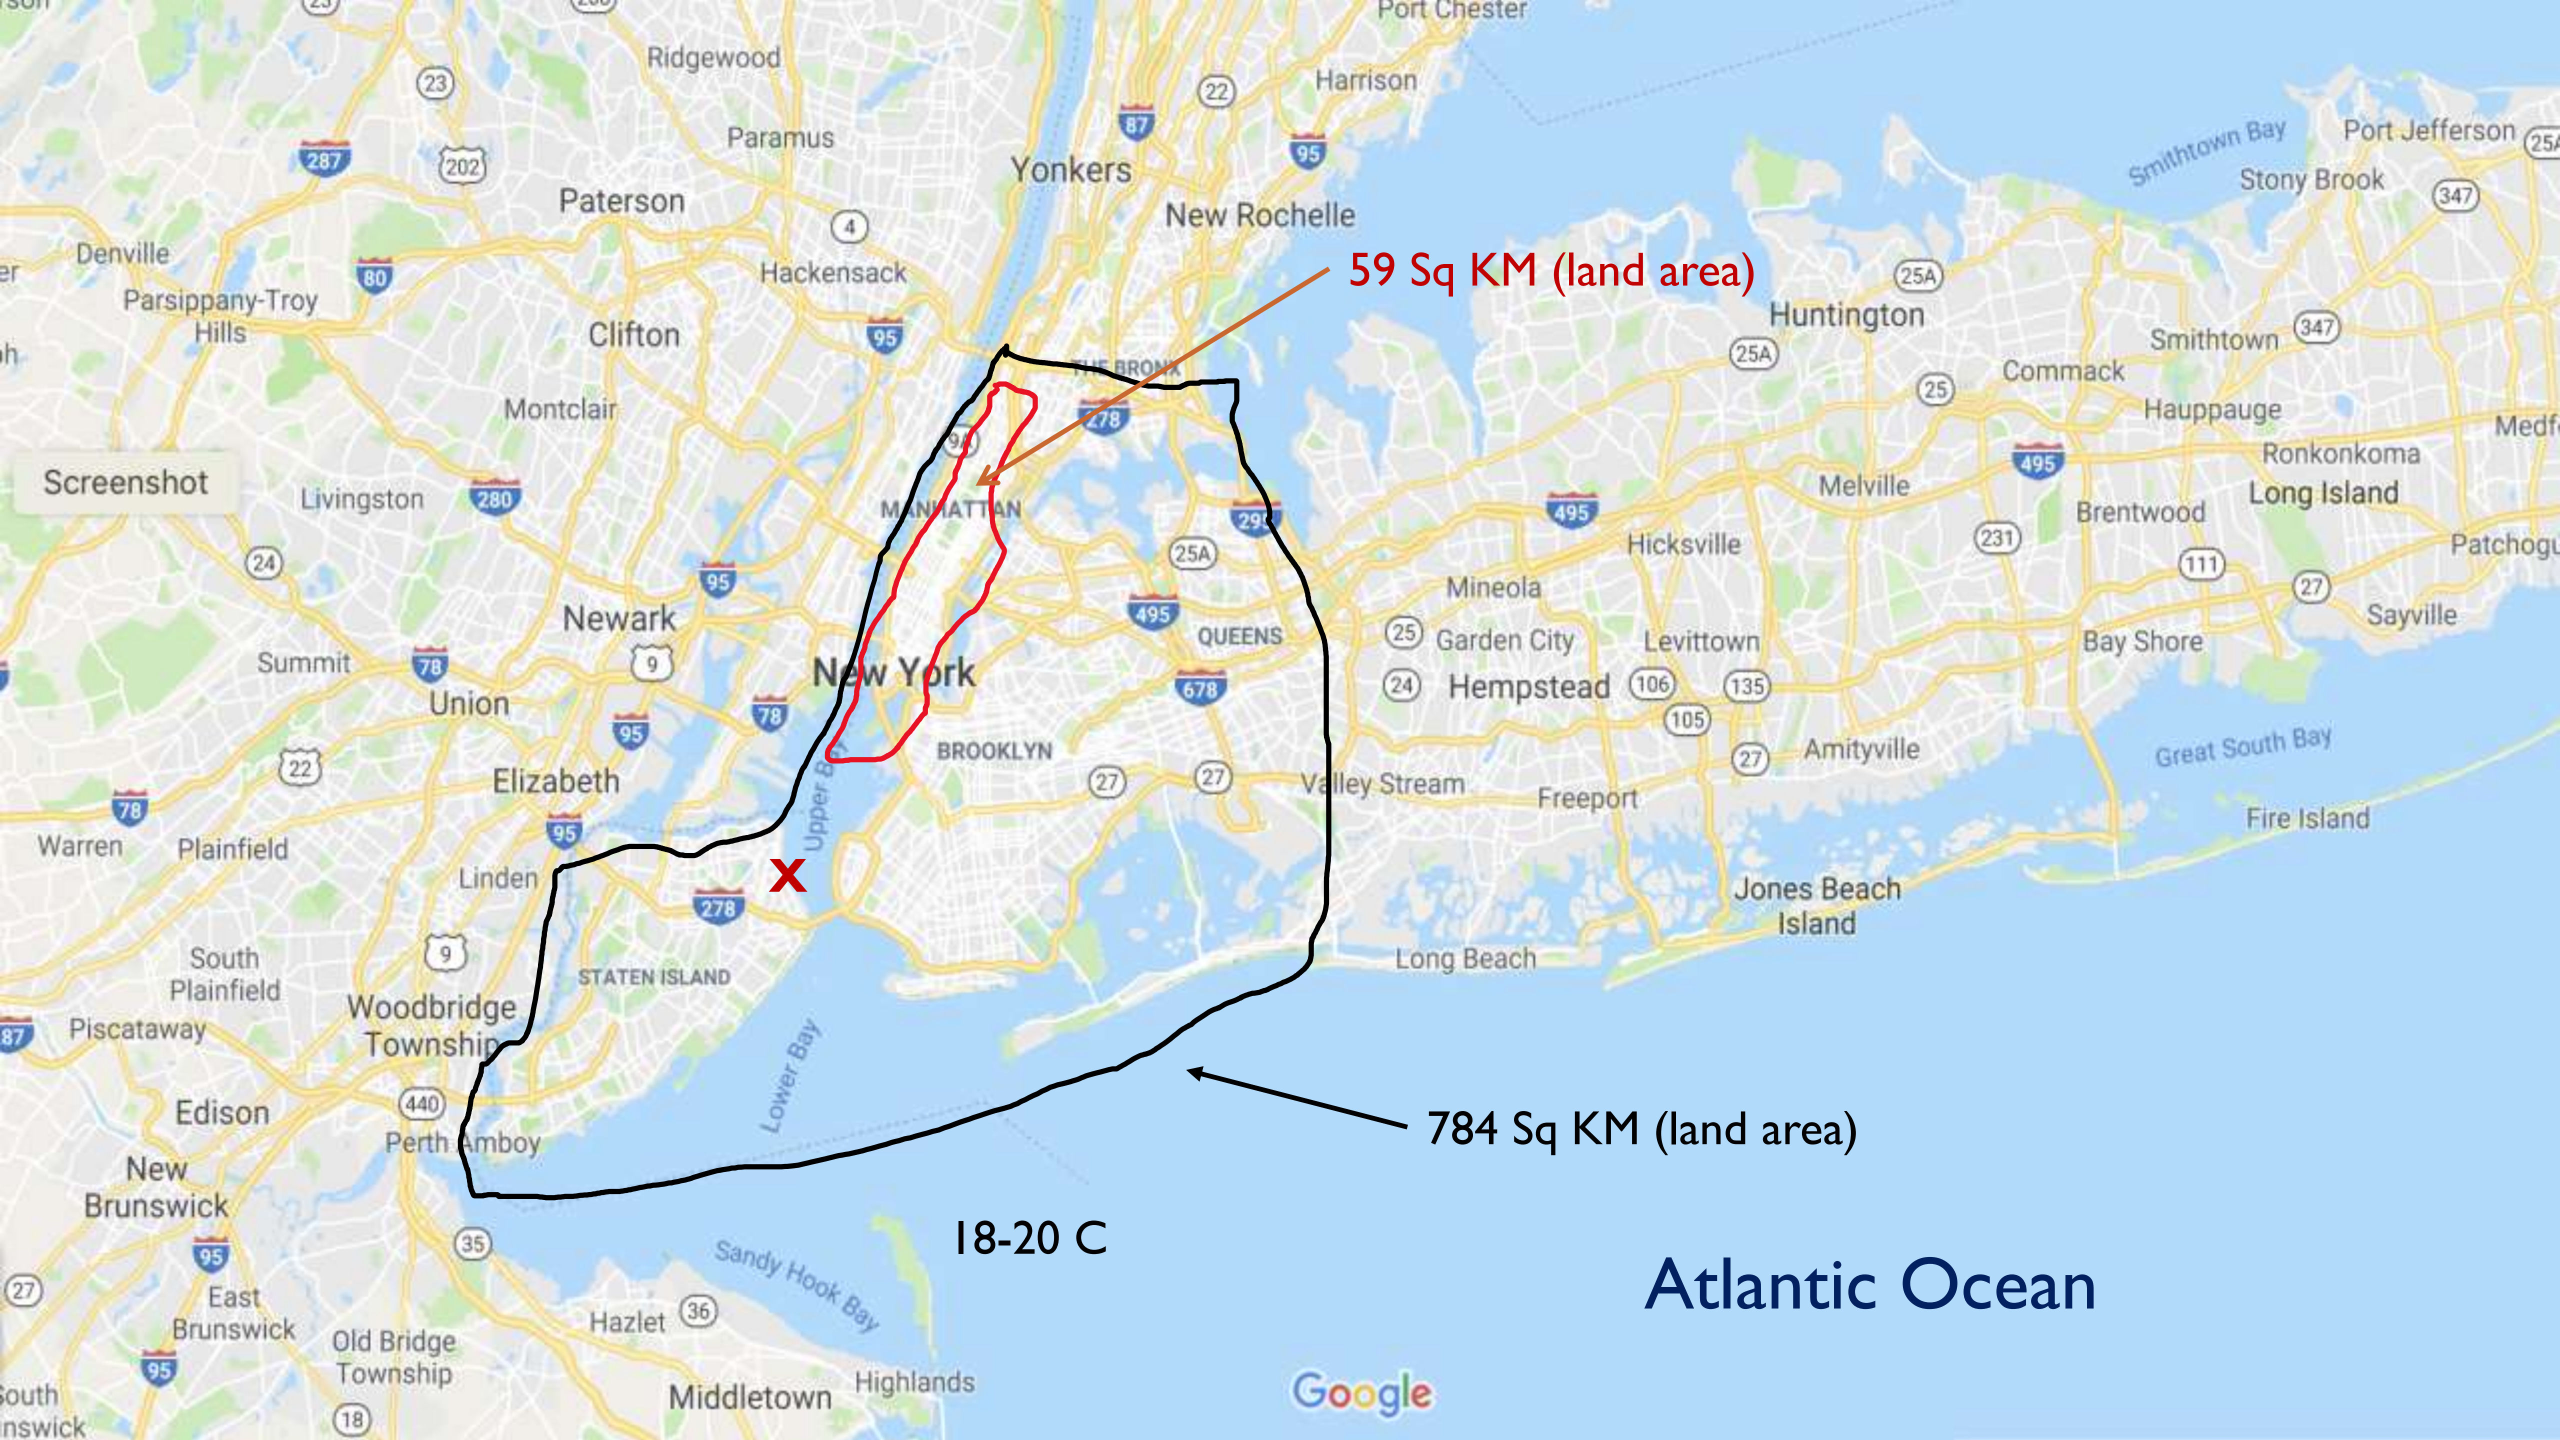
59 Sq KM (land area)
784 Sq KM (land area)
18-20 C
Atlantic Ocean
MANHATTAN
New York
THE BRONX
QUEENS
BROOKLYN
STATEN ISLAND
Screenshot
Google

0730, 30 September 2014
From Staten Island looking NE towards Manhattan

0730, 30 September 2014
From Staten Island looking NE towards Manhattan

PRESENTATIONS ON METADATA AND FILE SYSTEMS

- **Cybersecurity and Risk Management and World-wide Standards**
  - Henry Newman, CTO Seagate Government Solutions
- **Evolving NASA's Data and Information Systems for Earth Science**
  - Rahul Ramachandran, Senior Research Scientist | Earth Science Branch, NASA Marshall Space Flight Center
- **Requirements for Metadata and Archiving at Scale**
  - Amanda Tumminello, Navy Distributed Resource Center, DoD
- **iRODS: Metadata and Archiving at Scale**
  - Terrell Russell, Chief Technologist, iRODS Consortium
- **Moving from Extreme Scale Data to Extreme Scale Metadata Concerns: It's About Time!**
  - Gary Grider, HPC Division Leader, LANL/US DOE
- **Data Movement & Tiering with DMF 7**
  - Kirill Malkin, Director of Engineering, HPE
- **BeeGFS**
  - Frank Herold, CEO, ThinkParQ
- **PANASAS: Metadata Is Not Data**
  - Curtis Anderson, Panasas
- **MetaOcean: Handling Massive and Sensitive Metadata from Precision Medicine**
  - Frank Lee, IBM

https://hpcuserforum.com/presentations.html


PRICE $8.99
THE
NEW YORKER
SEPT. 30, 2019

CYBERSECURITY AND RISK MANAGEMENT AND WORLD-WIDE STANDARDS

HENRY NEWMAN, CTO SEAGATE GOVERNMENT SOLUTIONS

Digital Disruption

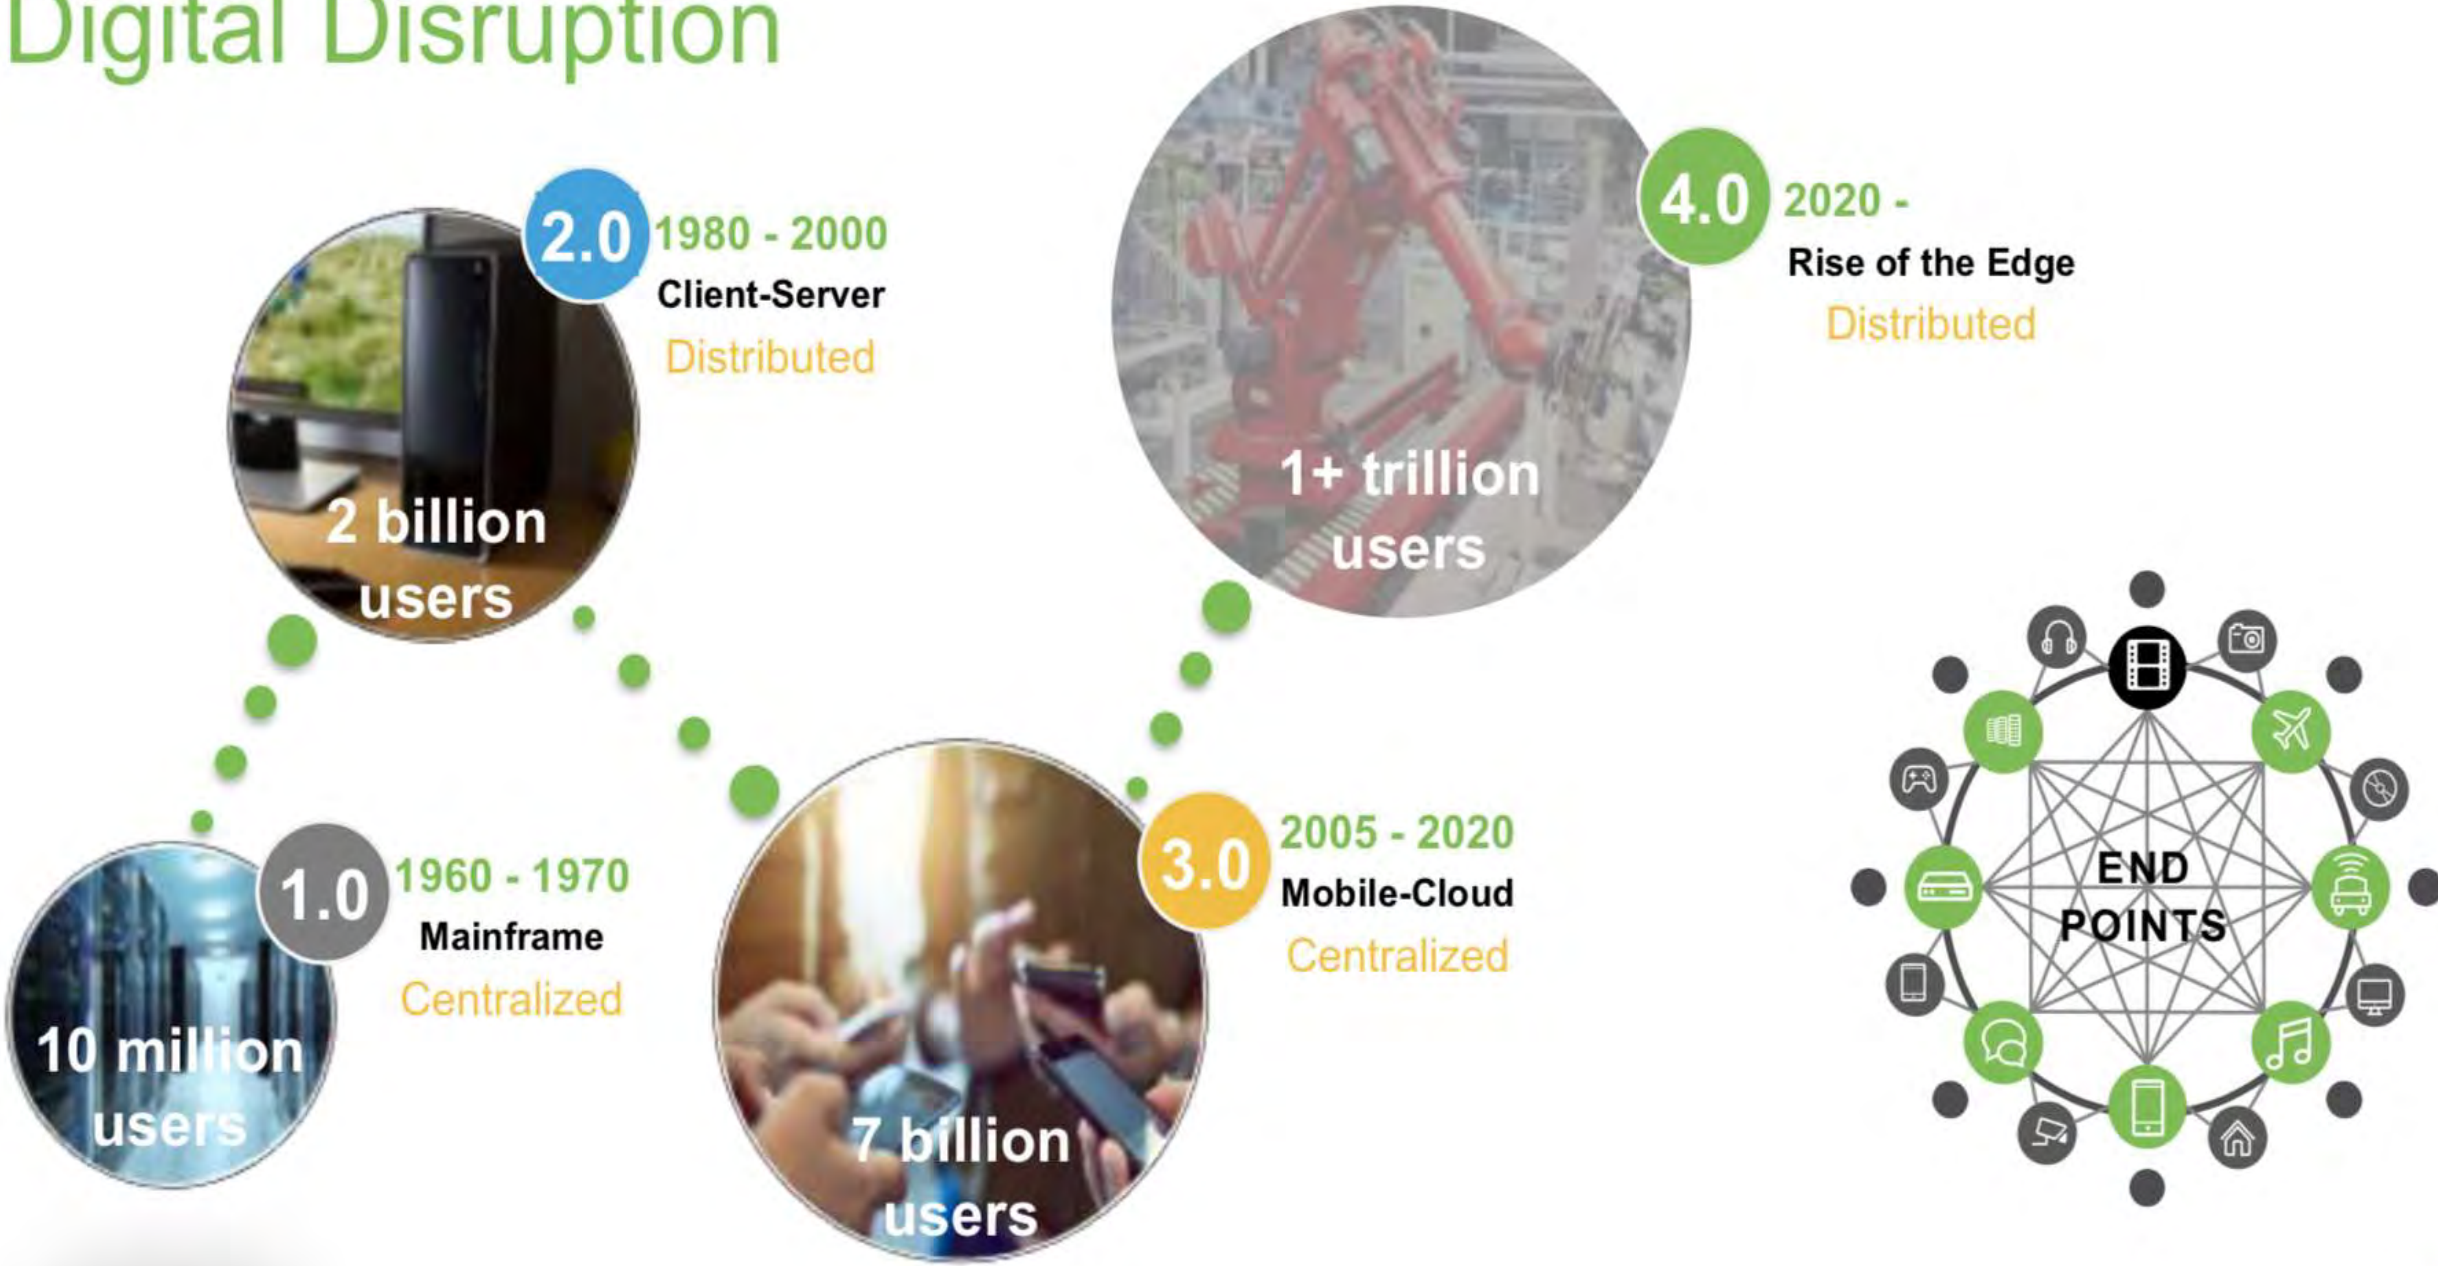

Security Certification and Standards

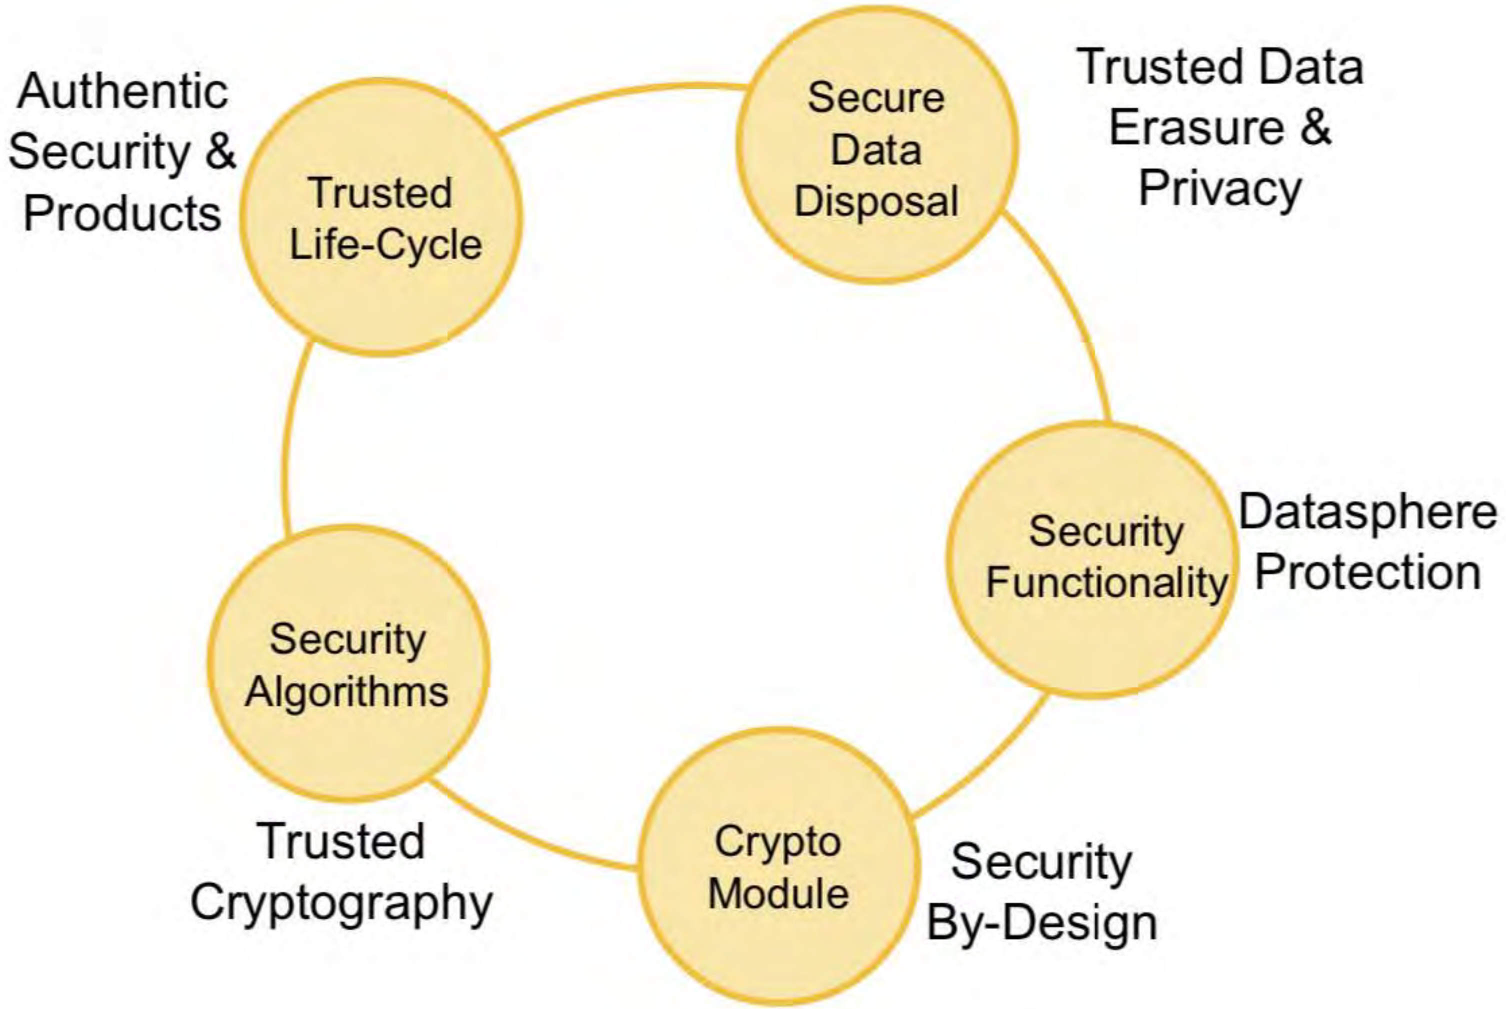

Security Algorithm Certifications

Security Algorithms

Trusted Cryptography

- Standard and Trusted Security Algorithms
- Certifications of all algorithms
  - Data Encryption
  - Integrity & Signatures
  - Random # Generation
  - Key Derivation…
- Required for FIPS 140-2 & Common Criteria Certs

NIST National Institute of Standards and Technology

Cryptographic Algorithm Validation Program (CAVP)

Security Module Certifications: FIPS 140-2

Crypto Module

Security By-Design

- Fundamental Security Certification
- Evaluation by Independent Labs
- Required for Information Security Products in Sensitive and Unclassified space in US & Canada
- Value recognized in other geographies

NIST National Institute of Standards and Technology

Communications Security Establishment Canada

Cryptographic Module Validation Program (CMVP)

Security Module Certifications: Common Criteria (CC)

Security Functionality

Datasphere Protection

- Security Use-Case (Protection Profile) Certification
- Evaluation by Independent Labs
- Certification recognized by 28 member nations globally for Information Security acquisition

Common Criteria for Information Security Evaluation (CC)

Sanitization Standard

Secure Data Disposal

Trusted Data Disposal & Privacy

- NIST SP 800-88 (Federal) & ISO 27040 (International) define media sanitization
- NIST SP 800-57 Defines Crypto Algorithm Longevity for erasure assurance.

NIST National Institute of Standards and Technology

ISO

[NIST Special Pub 800-88]

[NIST Special Pub 800-57]

ISO 27040

Crypto Algorithm Longevity*

Security Strength		2011 through 2013	2014 through 2030	2031 and Beyond
80	Applying	Deprecated	Disallowed	
	Processing	Legacy use		
112	Applying	Acceptable	Acceptable	Disallowed
	Processing			Legacy use
128	Applying/Processing	Acceptable	Acceptable	Acceptable
192		Acceptable	Acceptable	Acceptable
256		Acceptable	Acceptable	Acceptable

AES in any key size (128, 192, 256) is acceptable for use to 2031 and Beyond.

Trusted Life-Cycle Standards

THE Open GROUP

ISO 20243

Trusted Tech Provider Standard

Category	Section	Subsection
Technology Development	Product Development / Engineering Method	Software / Firmware / Hardware Design Process
		Configuration Management
		Well-Defined Development / Engineering Method Process and Practices
		Quality and Test Management
		Product Sustainment Management
	Secure Development / Engineering Method	Threat Analysis and Mitigation
		Run-time Protection Techniques
		Vulnerability Analysis and Response
		Product Patching and Remediation
		Secure Engineering Practices
		Monitor and Assess the Impact of Changes in the Threat Landscape
Supply Chain	Supply Chain Security	Risk Management
		Physical Security
		Access Controls
		Employee and Supplier Security and Integrity
		Business Partner Security
		Supply Chain Security Training
		Information Systems Security
		Trusted Technology Components
		Secure Transmission and Handling
		Open Source Handling
		Counterfeit Mitigation
		Malware Detection

Trusted Life-Cycle

Authentic Security & Products

- The Open Trusted Technology Provider Standard (O-TTPS) is now a sanctioned ISO Standard
- Comprehensive Secure Technology Provider Standard
- Sections for Secure Technology Development and Secure Supply Chain
- The NIST Cybersecurity Framework Provides for common framework and language for managing Cyber Risk

NIST National Institute of Standards and Technology

Cybersecurity Framework

EVOLVING NASA’S DATA AND INFORMATION SYSTEMS FOR EARTH SCIENCE

RAHUL RAMACHANDRAN, SENIOR RESEARCH SCIENTIST | EARTH SCIENCE BRANCH, NASA MARSHALL SPACE FLIGHT CENTER

Extensive Data Collection

Started in the 1990s, EOSDIS today has 11,000+ data types (collections)

Land
- Cover & Usage
- Surface temperature
- Soil moisture
- Surface topography

Ocean
- Surface temperature
- Surface wind fields & heat flux
- Surface topography
- Ocean color

Atmosphere
- Winds & Precipitation
- Aerosols & Clouds
- Temperature & Humidity
- Solar radiation

Human Dimensions
- Population & Land Use
- Human & Environmental Health
- Ecosystems

Cryosphere
- Sea/Land Ice & Snow Cover

Rahul Ramachandran

Earth Observing System Data and Information System (EOSDIS)

EOSDIS is managed by the Earth Science Data and Information System (ESDIS) Project at GSFC and includes the following major core components:

<u>Science Investigator-led Processing Systems (SIPS)</u>

- Perform forward processing of standard data products and reprocess data to incorporate algorithm improvements

<u>Distributed Active Archive Centers (DAACs)</u>

- Co-located with centers of science discipline expertise; archive and distribute standard data products produced by the SIPS and others

<u>Earthdata and Core Services</u>

- Allows users to search, discover, visualize, refine, and access NASA Earth Observation data. Includes networking and security

Red highlight indicates EOSDIS boundary.

Rahul Ramachandran

Earth Observing System Data and Information System (EOSDIS)

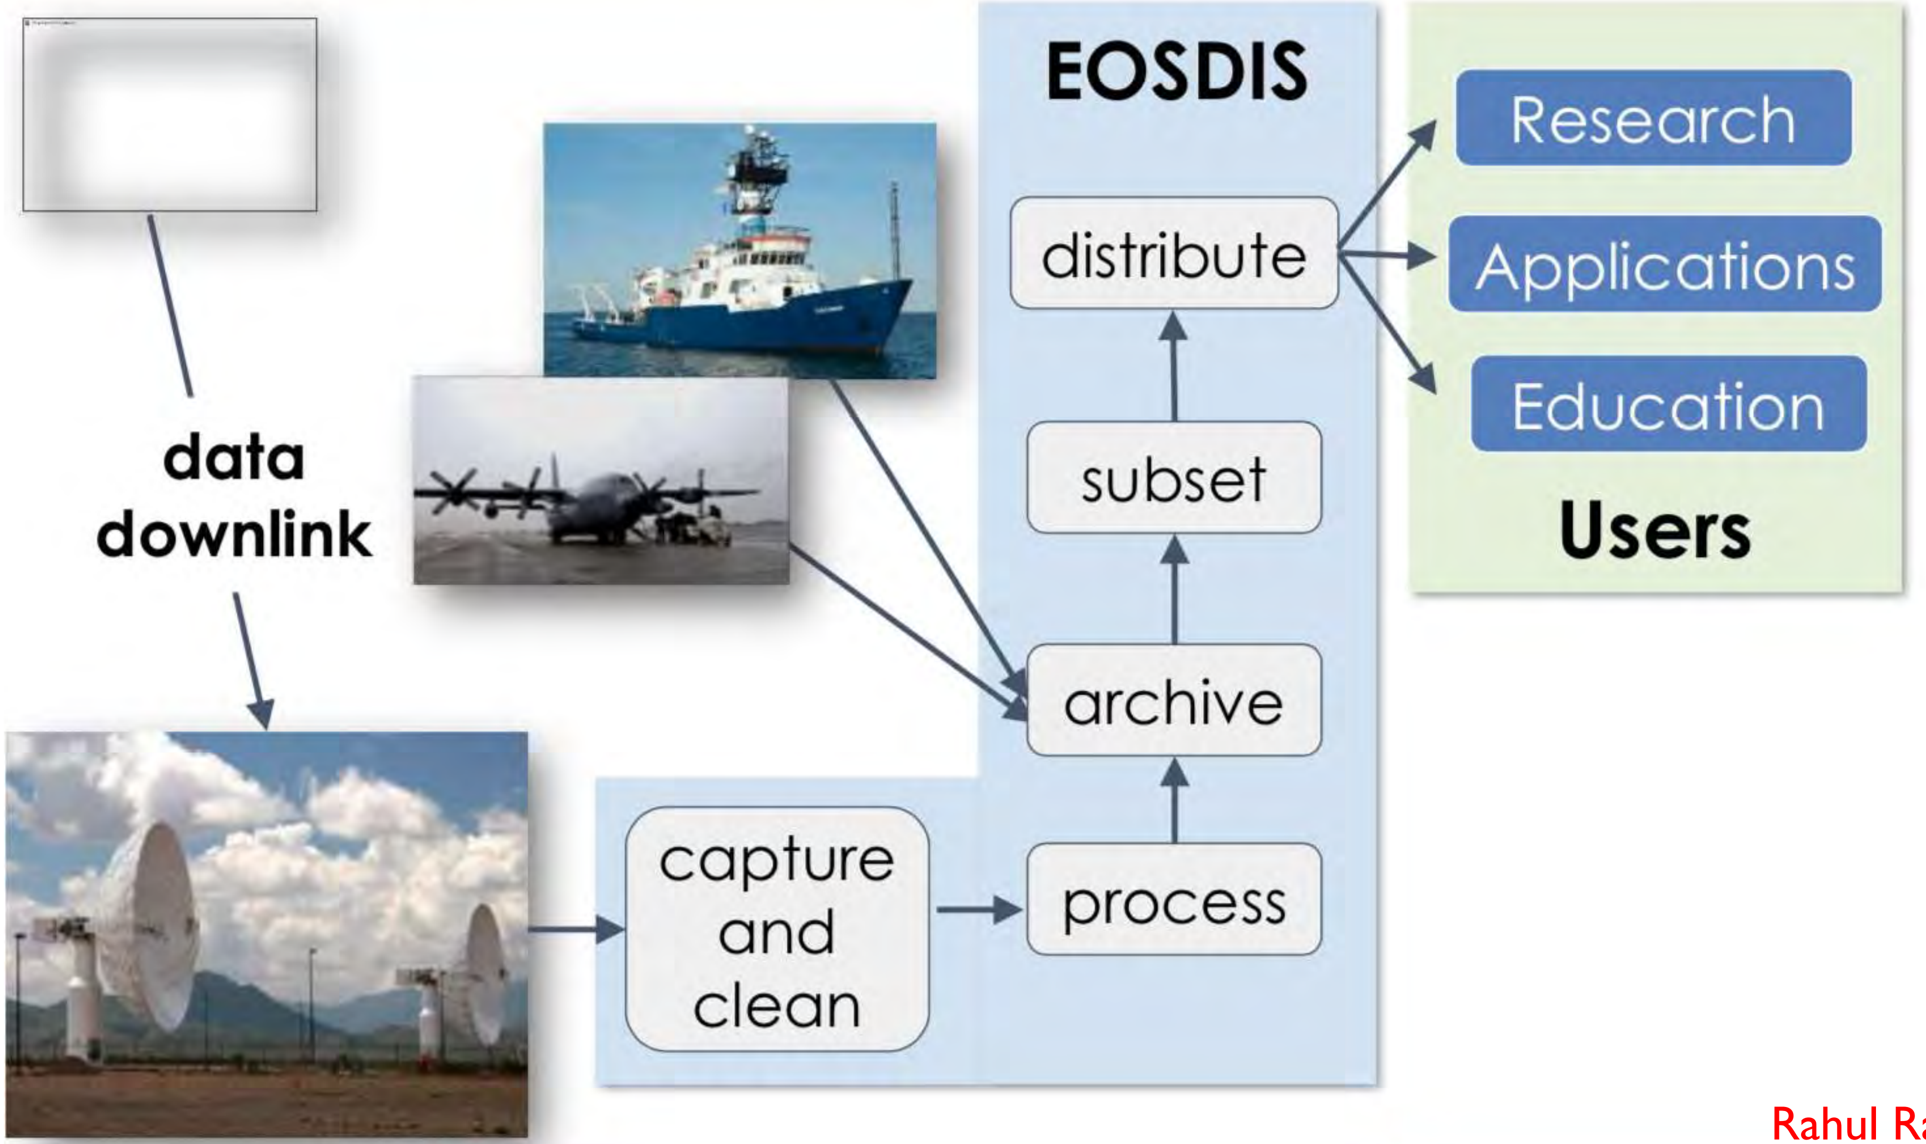

Rahul Ramachandran

Data Sources

Type	Example Missions
Satellite/on-orbit Missions	Terra, Aqua, Aura, Suomi-NPP, SORCE, GPM, GRACE, CloudSat, CALIPSO, etc.
Airborne Missions	IceBridge, Earth Ventures (5+ missions), UAVSAR, etc.
In Situ Measurement Missions	Field campaigns on land (e.g., LBA-ECO) and in the ocean (e.g., SPURS)
Applications support	Near-real time creation and distribution of selected products for applications communities
Earth Science Research support	Research products from efforts like MEaSUREs. This also includes data from older, heritage missions (prior to EOS Program) that the DAACs rescue – e.g., Nimbus, SeaSat

Rahul Ramachandran

Role of DAACs

DAACs were selected and established based on the Earth Science discipline expertise and heritage of their host organizations.

- Provide unique support and expert services to their user communities
- Provide data and services to the research community for comprehensive, cross-discipline studies needed to understand Earth as an interrelated system
- Ensure safe stewardship of NASA's data

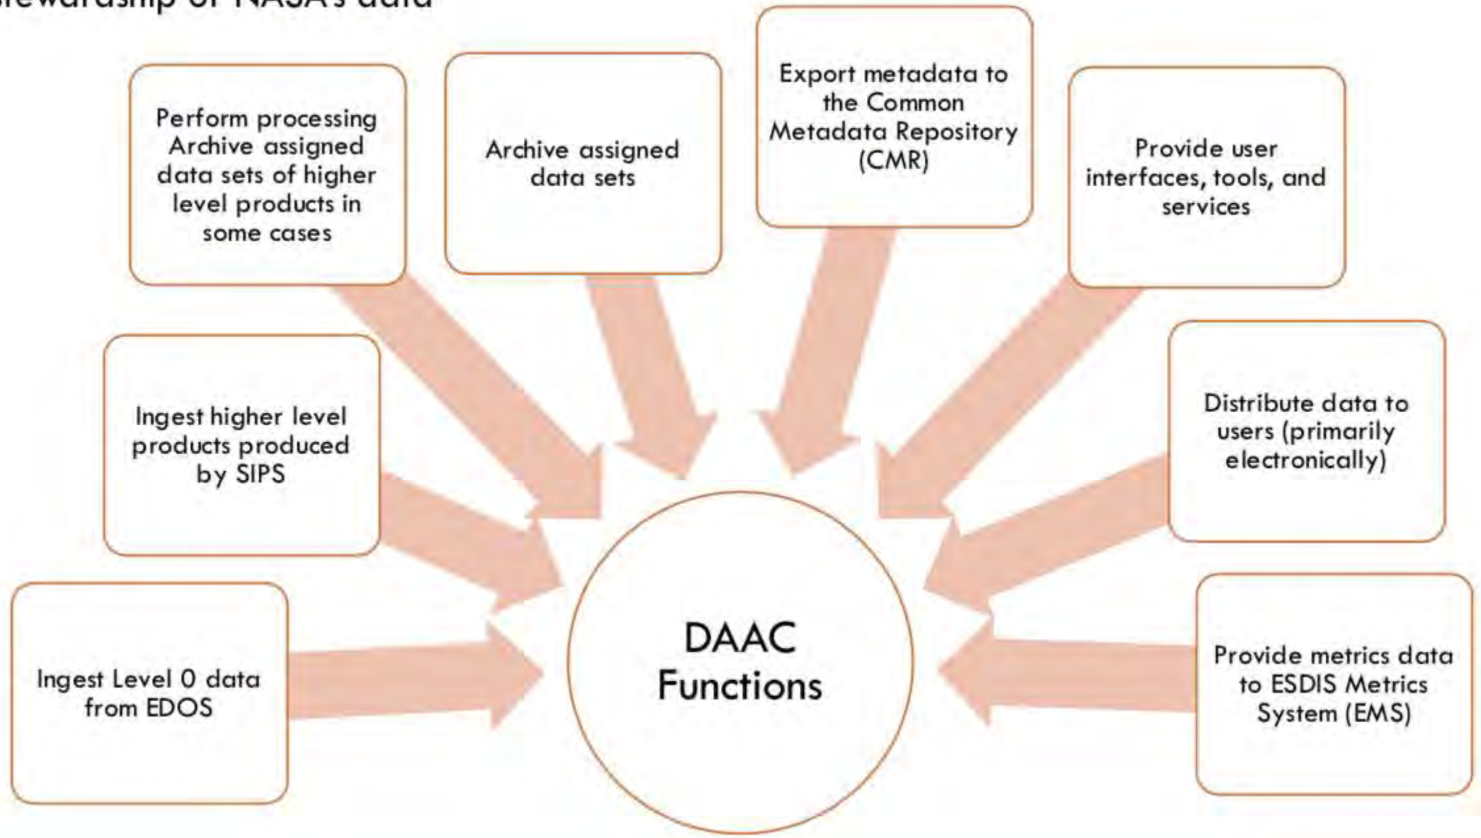

Rahul Ramachandran

Procedures for Archiving

No matter what type of data (on orbit, aircraft, in situ, etc.) DAAC Staff follow these procedures for handling datasets

Planning	• Collaborate with mission teams, data producers, and ESDIS to develop Interproject Agreements (IPAs), Interface Control Documents (ICDs), and Operations Agreements (OAs) • Support data producers with the creation of Data Management Plans (DMPs)
Acquire and/or Produce Data	• Advise data producers on data formats, structure, and delivery methods • Establish automated processes for the transfer of data into DAAC data systems • Develop or integrate, test, verify, and run data production code (for applicable data)
Preserve Data	• Ensure redundant online disk and tape data storage • Establish file management (e.g. duplicate file detection) and file integrity (e.g. checksum verification)
Describe Data	• Create collection-level metadata and, when necessary, employ software to extract file-level metadata for the purposes of preservation, discovery and usage • Export metadata to NASA's CMR for inter-mission and inter-sensor data discovery • Develop user documentation and supplemental information • Create DOIs and data citations for proper attribution of data and data creators
Distribute Data	• Provide discovery through the DAAC Web site, with direct HTTPS access • Support automated data transfer to users through subscriptions and APIs • Facilitate search, visualization, and customization through NASA Earthdata Search • Develop specialized portals and data services per mission and user needs
Support Data	• Assist user communities with the selection and usage of data • Create "How To" guides, FAQs, and other resources to address specific user needs • Work with user communities to identify needed improvements for data and tools • Provide outreach and education to broaden the user community

Core Services

Metrics!

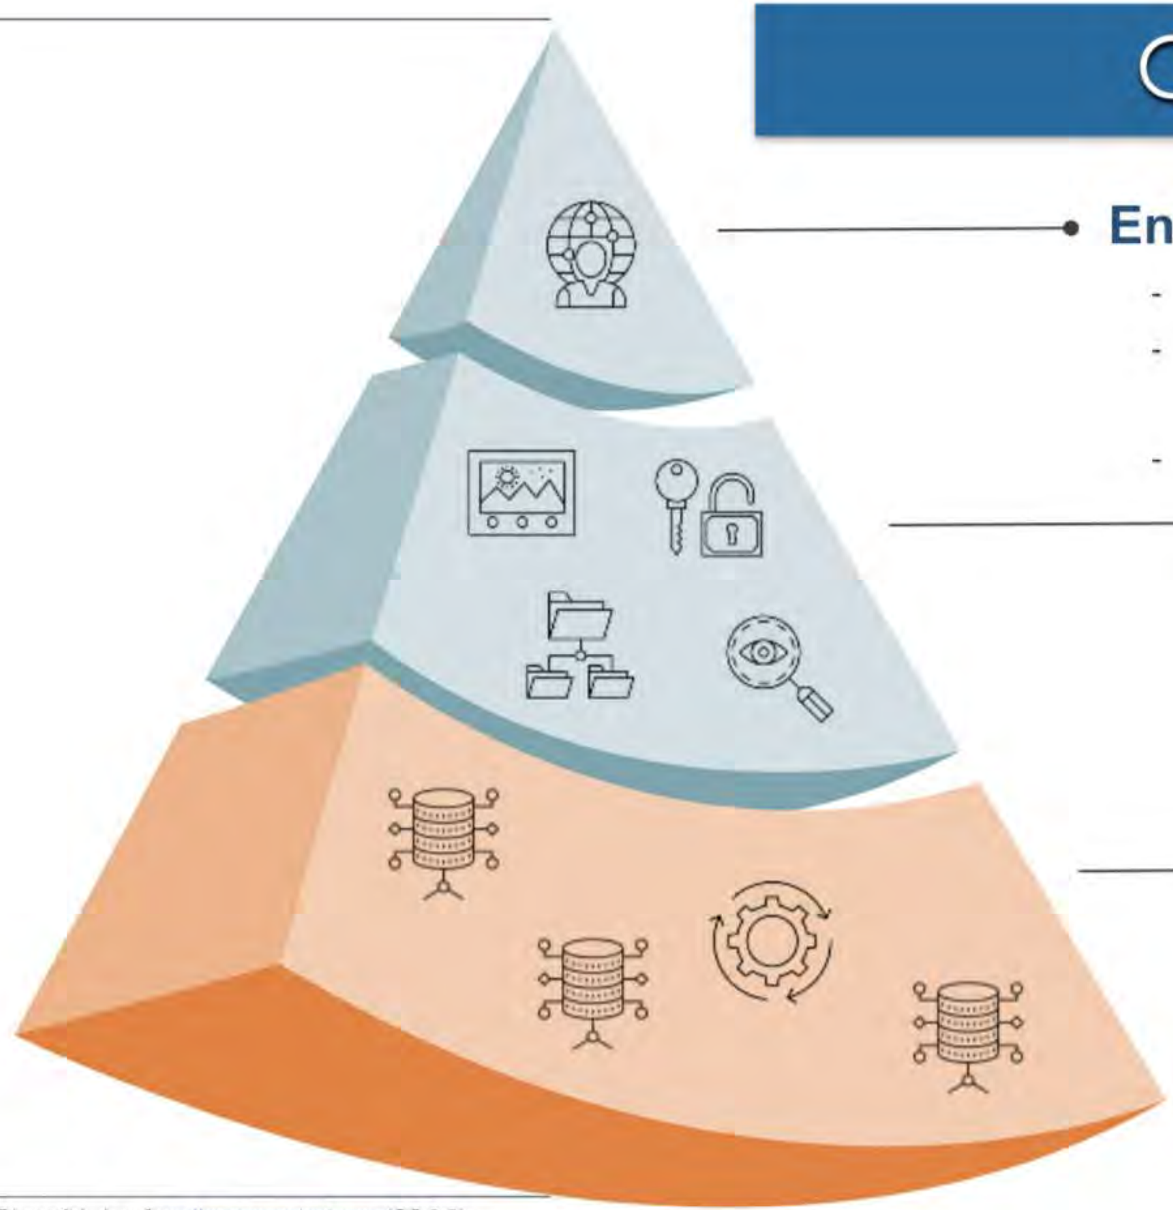

all by Dinosoft Labes from thenounproject.com (CC 3.0)

- **End User Web Clients**
  - **Earthdata.nasa.gov**
  - **Earthdata Search:** data access/discovery*
  - **Worldview:** imagery*
- **Open Service APIs**
  - **CMR**: Metadata Catalog (Search Engine)*
  - User Login
  - **GIBS**: Global Imagery Browse Services*
- **Earth Science Data Holdings**
  - Open APIs
  - Free Data Download
  - DAAC specific tools

EOSDIS core tools

Federated resources

*open source software

Common Metadata Repository

Provides a single source of unified, high-quality, and reliable Earth Science metadata with a high performance ingest and search architecture for submission and discovery of all EOSDIS data sets.

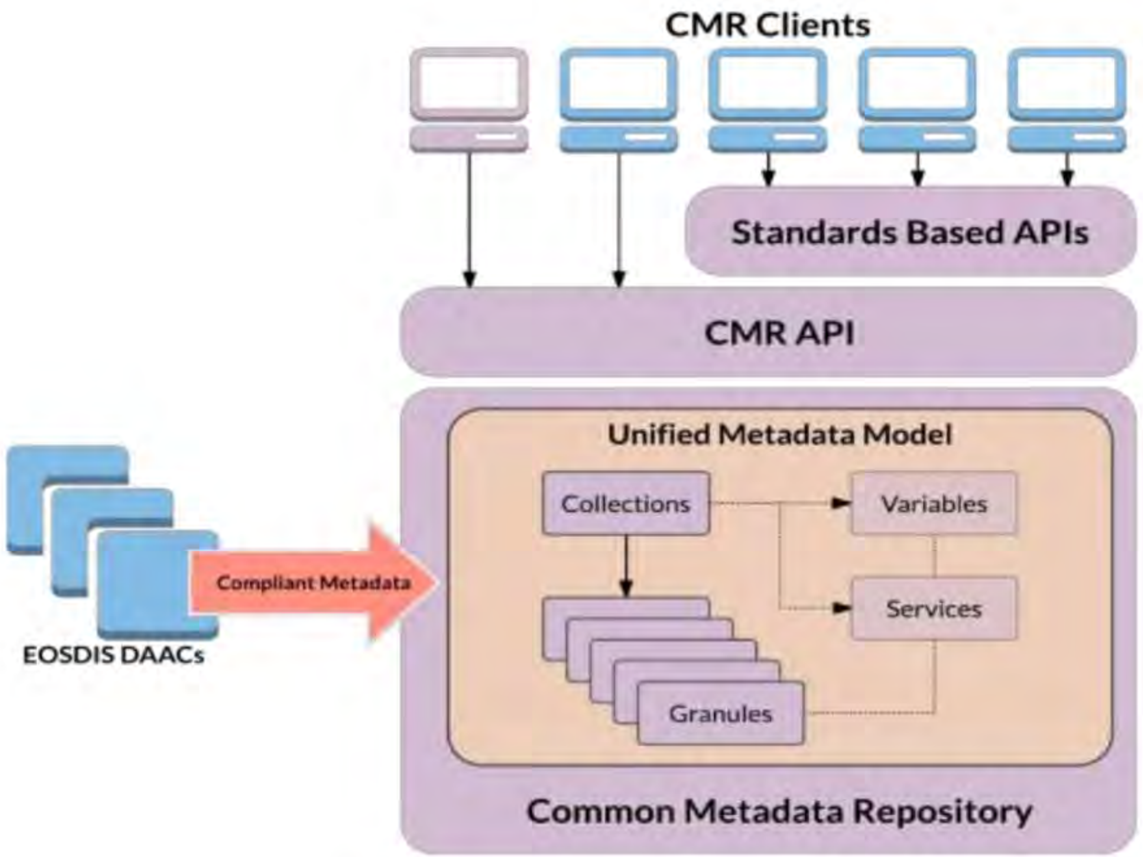

Lightning fast, always available

- *95% queries complete in <1s*
- *99.98% uptime (last 365d)*

Big Data Ready

- *34K collections*
- *367 million files indexed*
- *Prepared to scale 1B+ records*

Standards-focused

- *ISO-19115 metadata*
- *OpenSearch/OGC CSW*
- *REST based APIs*

Community-focused

- *Developer's portal*
- *Active Developer's forum*
- *Ecosystem of supported tools*
- *Open Source codebase in github*

Internationally Recognized

- *Provides the backbone of the Community of Earth Observing Satellites International Directory Network (CEOS IDN)*

Rahul Ramachandran

Common Metadata Repository (CMR) APIs

- CMR - spatial and temporal metadata registry
- Based on Extensible Markup Language (XML) and web service technologies
- Interfaces with clients and users through various APIs
- Open system
    - Access Control API Overview
    - API Documentation
    - Client Developer Forum
    - Client Partner User Guide
    - Collection CSW
    - Harvesting Best Practices
    - Metadata Ingest API Overview
    - Metadata Management Tool User Guide
    - OpenSearch Documentation
    - UMM C-Schema and S-Schema

REQUIREMENTS FOR METADATA AND ARCHIVING AT SCALE

Amanda Tumminello, US Navy HPC Center

What a user wants

- **Metadata**
  - Ability to find physical location of all copies of data.
  - How many files and total capacity being used (how much is left?)
  - I/O characteristics of the physical residence
  - Staging of data to specific resources
  - Chain of custody (user and digital access)
  - Extended attributes that are easily searchable
  - Ability to enhance/add/change attributes as environment and science change
- **Archive**
  - Ability to create transportable archives (physical/cloud)

Amanda Tumminello

What an admin needs

- **Real time I/O characteristics**
- **Easily queried sensitivity levels for data**
- **Encryption capabilities and key management**
  - Data at Rest Encryption
  - Data in Flight Encryption
  - Data Encryption per Object
  - User defined Encryption
  - Group defined Encryption

Screenshot

Amanda Tumminello

What and admin also needs

- Chain of Custody
- User, Project, Organizational level reporting for utilization and data location

Screenshot

- Data collection and movement for user/project/organization
- Data curation for lifetime of project
- Heat maps of metadata operations during processing
- Efficient means to purge data
- Ability to prove Qos of I/O subsystem

Amanda Tumminello

Balancing needs within an acquisition

- **Detailed reports (how/what/when/where)**
  - User level reports, System level reports (high level), Physical reports on data at rest (low level), Program level reports (enterprise level), Reports on data in flight (how is data moving)
- **Auditing of lifecycle of data**
- **Analysis of metadata types to data capacity for cost analysis**
- **Amount of time file/object spends on media type (SSD, HDD, tape)**
- **Heat maps of I/O traffic and utilization across the Enterprise**

Amanda Tumminello

IRODS: METADATA AND ARCHIVING AT SCALE

TERRELL RUSSELL, CHIEF TECHNOLOGIST, IRODS CONSORTIUM

Open Source Data Management

iRODS

Data Centric. Metadata Driven.

Provides insurance against your changing infrastructure:

- edge devices (sequencers, satellites, supercomputers, etc.)
- storage
- compute
- networking
- authentication

Terrell Russell

Why Data Management Matters

DATA LIFECYCLE

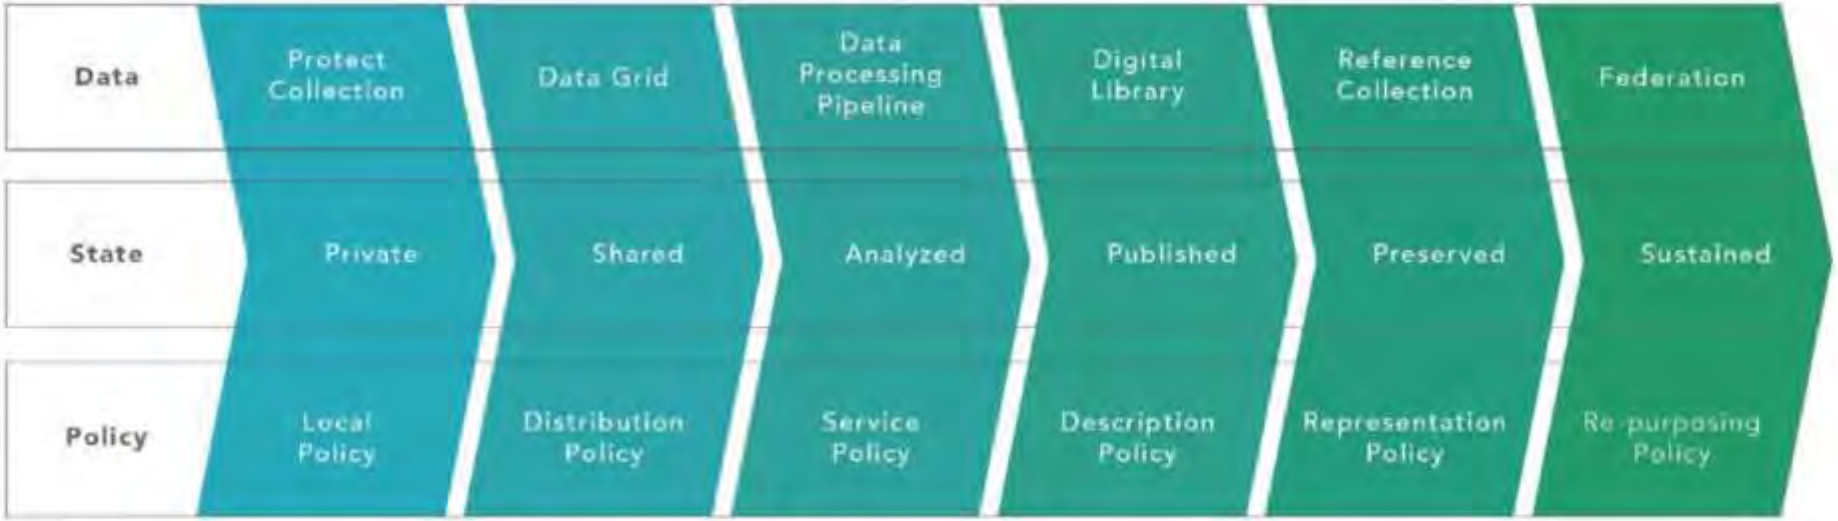

iRODS virtualizes the stages of the data lifecycle through policy evolution

As data matures and reaches a broader community, data management policy must also evolve to meet these additional requirements.

Terrell Russell

iRODS Core Competencies

iRODS

The underlying technology categorized into four areas

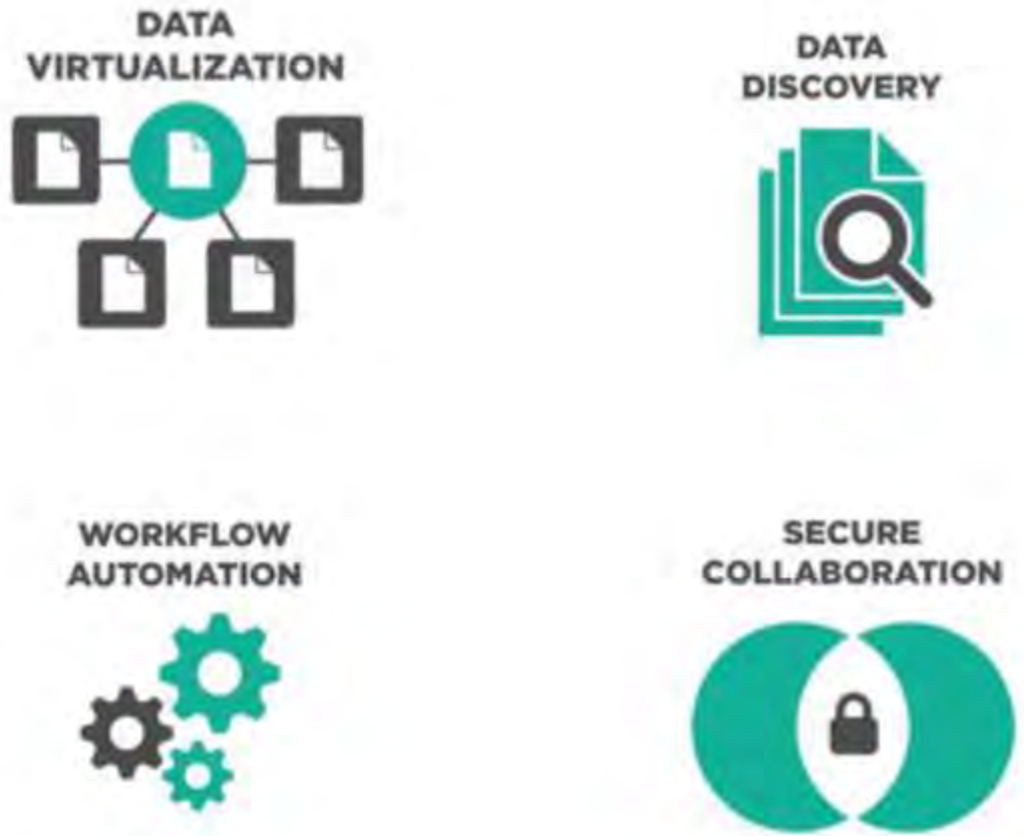

Terrell Russell

iRODS Capabilities

iRODS

- Automated Ingest
- Storage Tiering
- Auditing
- Provenance
- Indexing
- Publishing
- Data Integrity
- Compliance

Terrell Russell

Deployment Patterns

iRODS

Data to Compute

Compute to Data

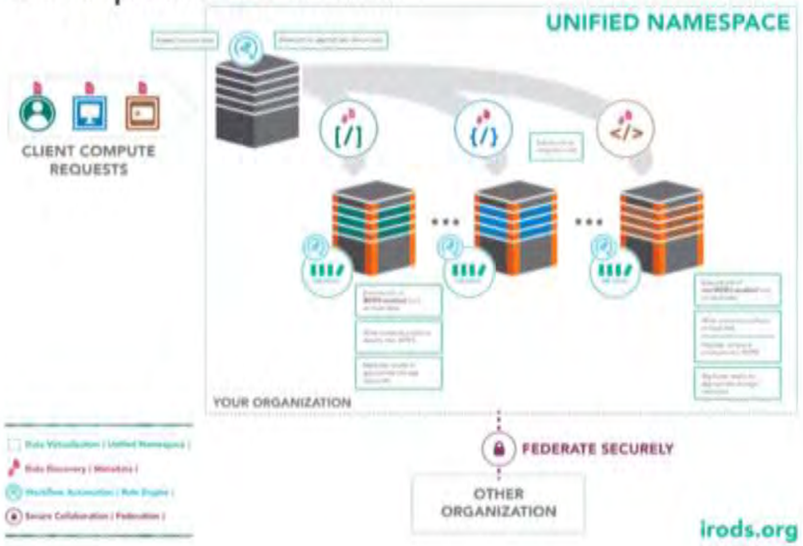

Filesystem Synchronization

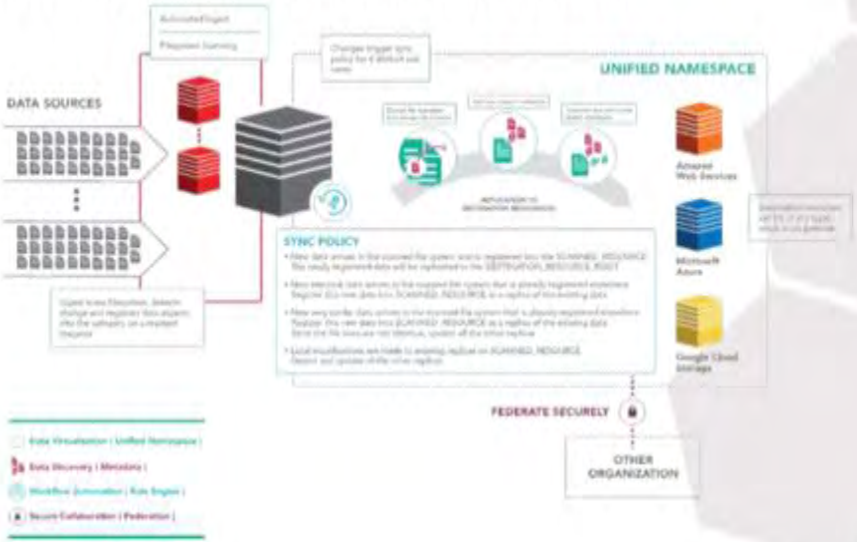

Terrell Russell

Deployment Patterns

iRODS

Data to Compute

Filesystem Synchronization

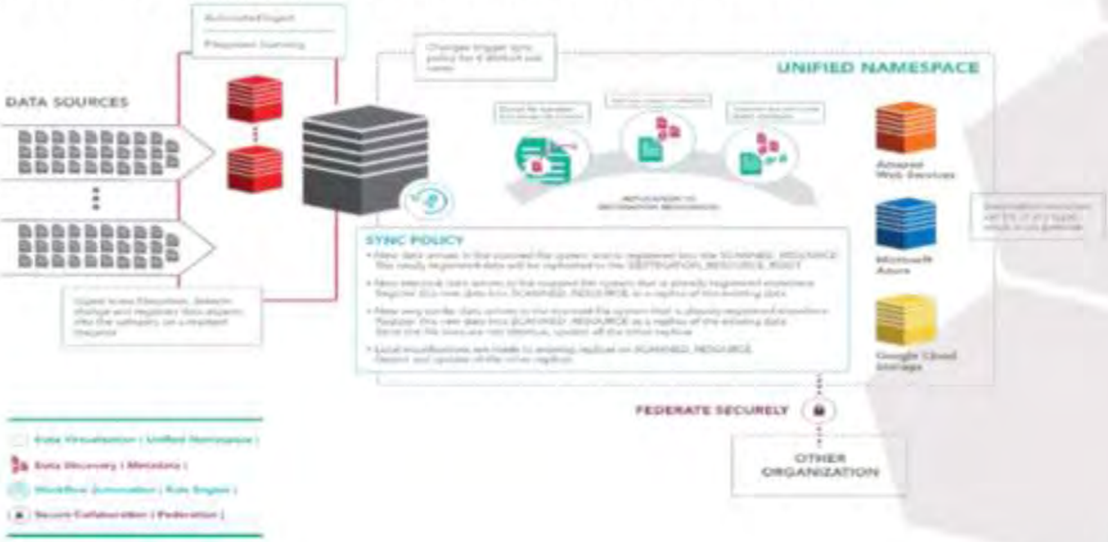

Compute to Data

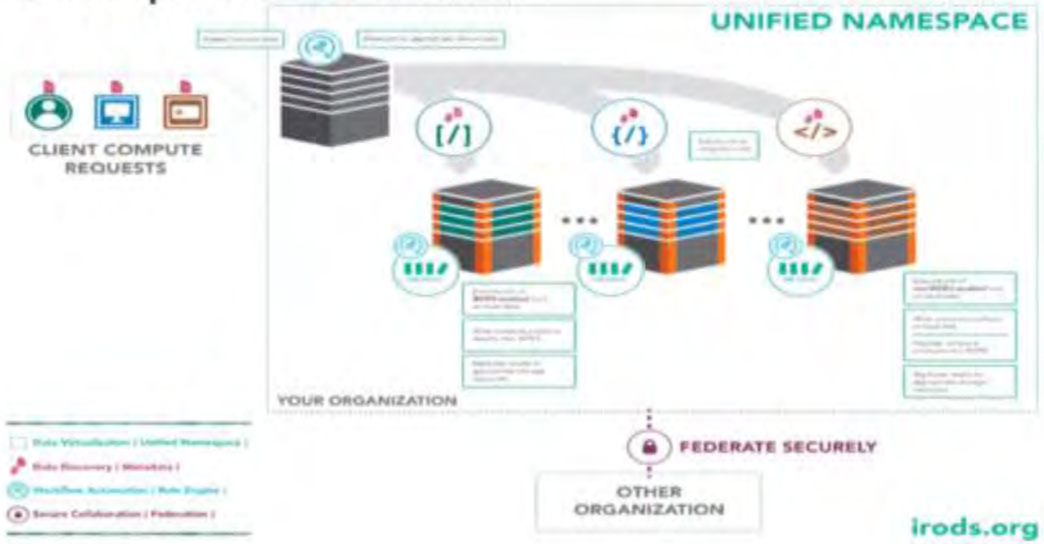

irods.org

Terrell Russell

The Data Management Model

iRODS

iRODS provides eight packaged capabilities, each of which can be selectively deployed and configured.

These capabilities represent the most common use cases as identified by community participation and reporting.

The flexibility provided by this model allows an organization to address its immediate use cases.

Additional capabilities may be deployed as any new requirements arise.

A pattern represents a combination of iRODS capabilities and data management policy consistent across multiple organizations.

Three common patterns of iRODS deployment have been observed within the community:

Automated Ingest - Landing Zone

iRODS

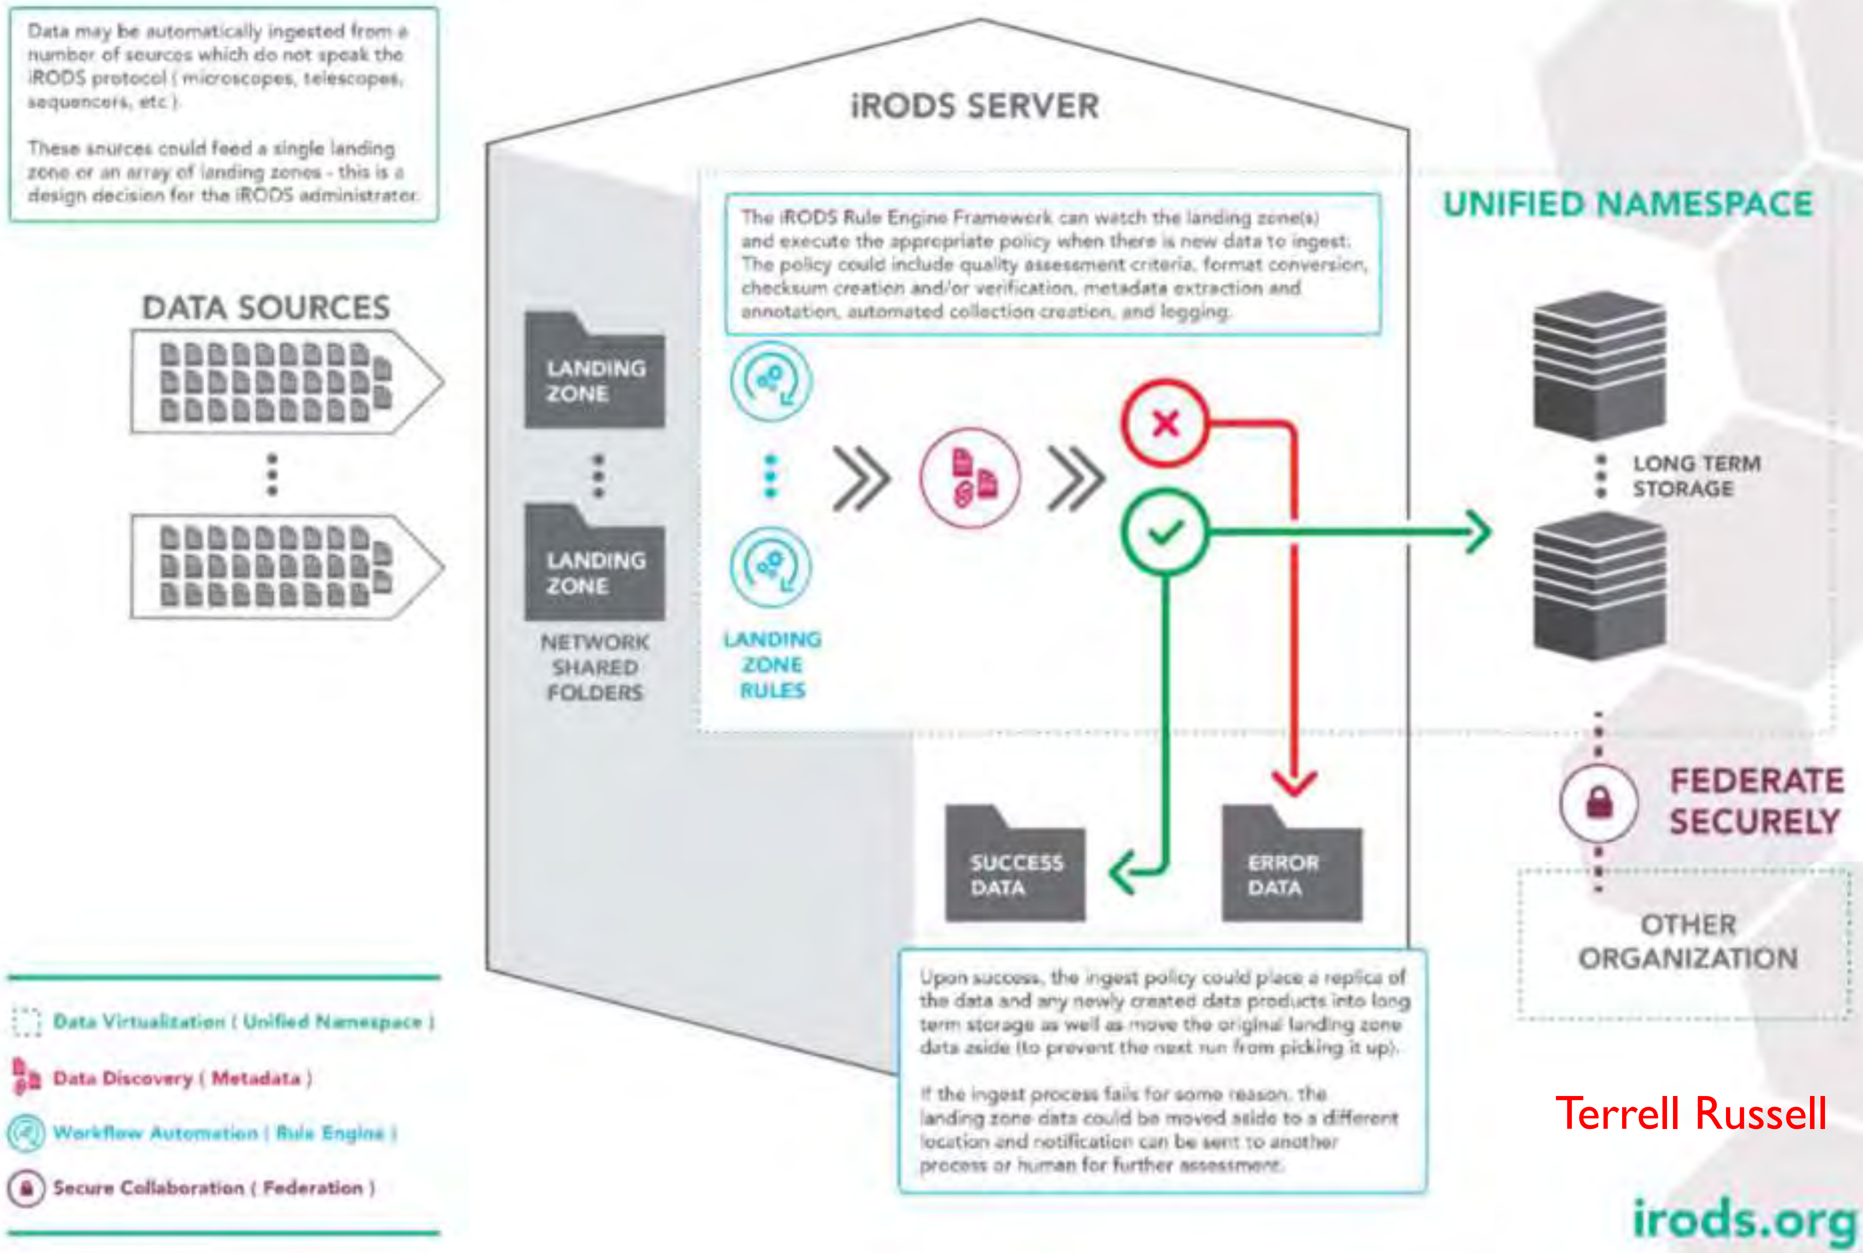

Terrell Russell

Automated Ingest - Filesystem Scanning

iRODS

Periodically a scanning job is added to the queue which generates jobs to register or ingest data.

Metadata is extracted and applied once the objects are registered in the catalog

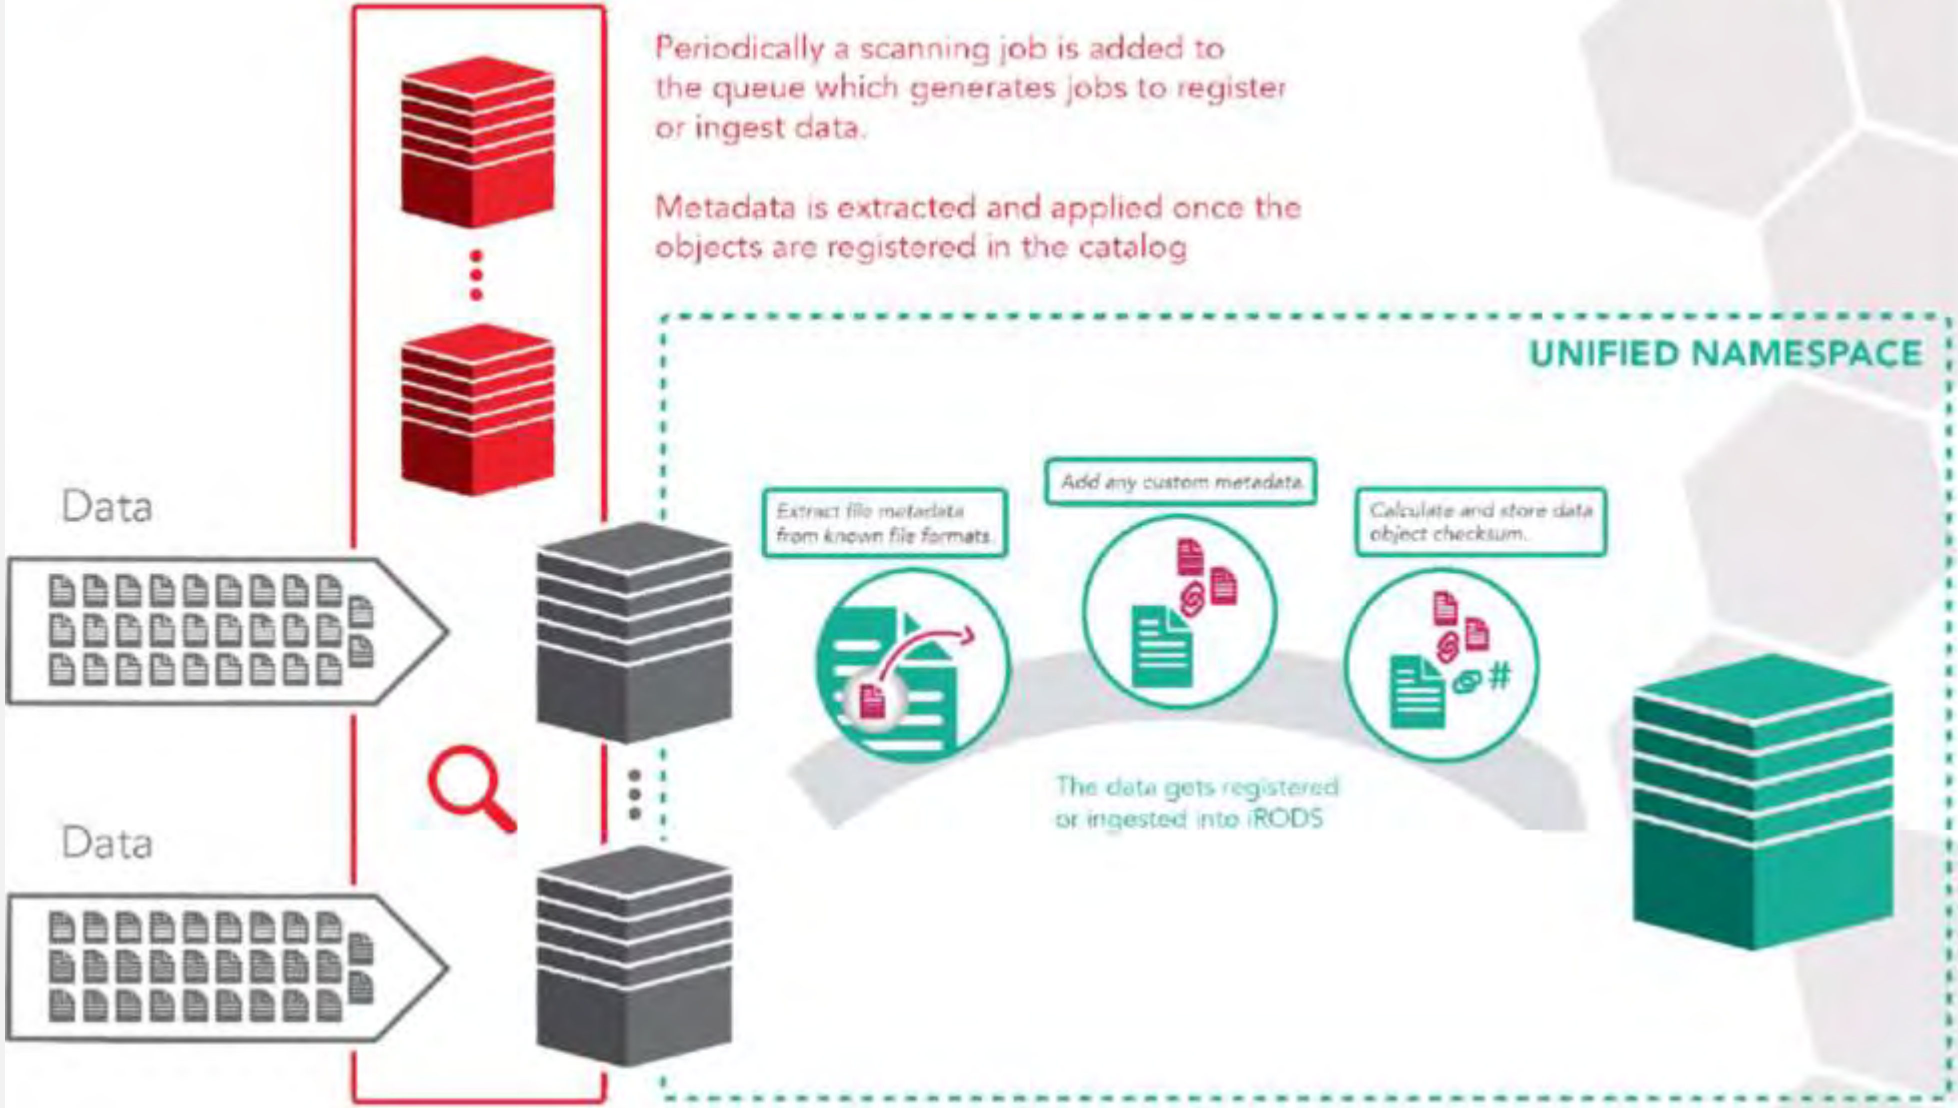

Terrell Russell

Automated Ingest - Filesystem Scanning

iRODS

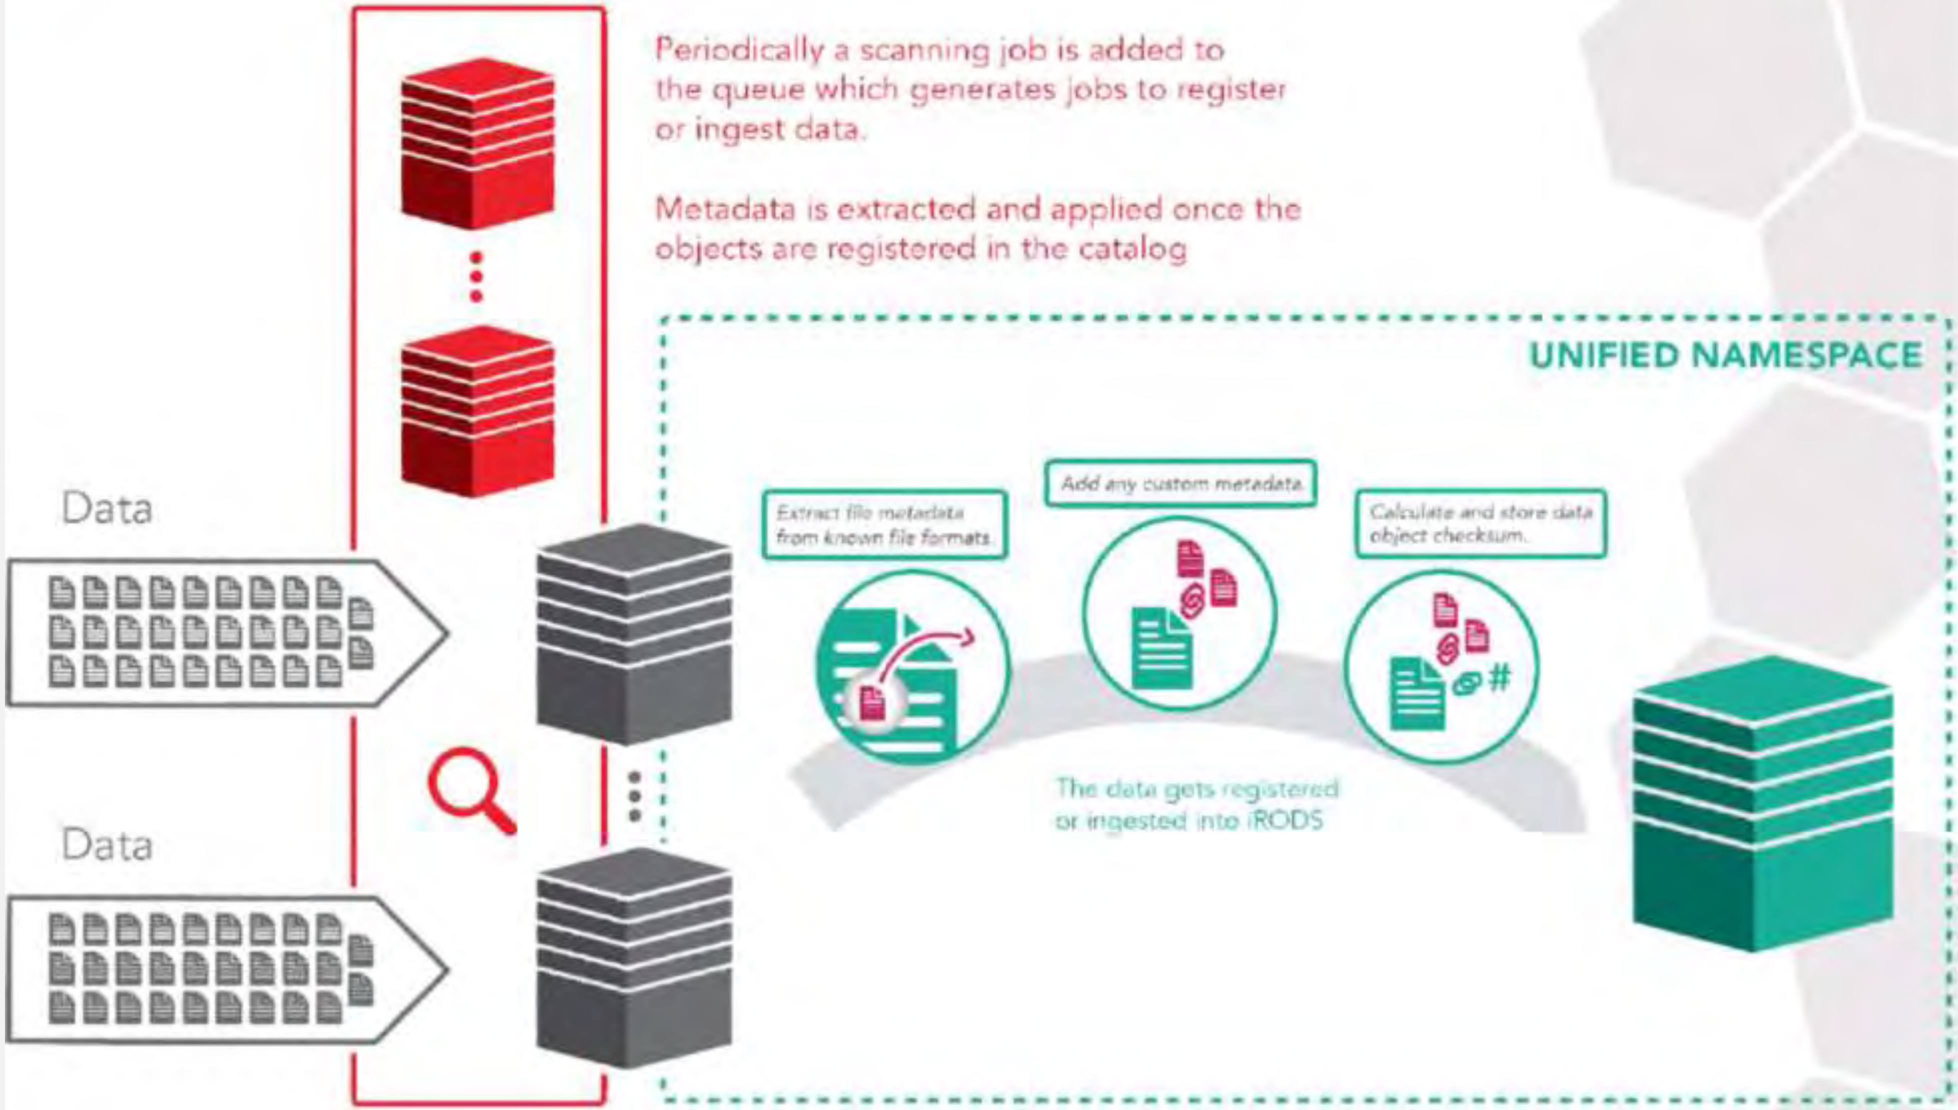

Terrell Russell

Compute to Data

iRODS

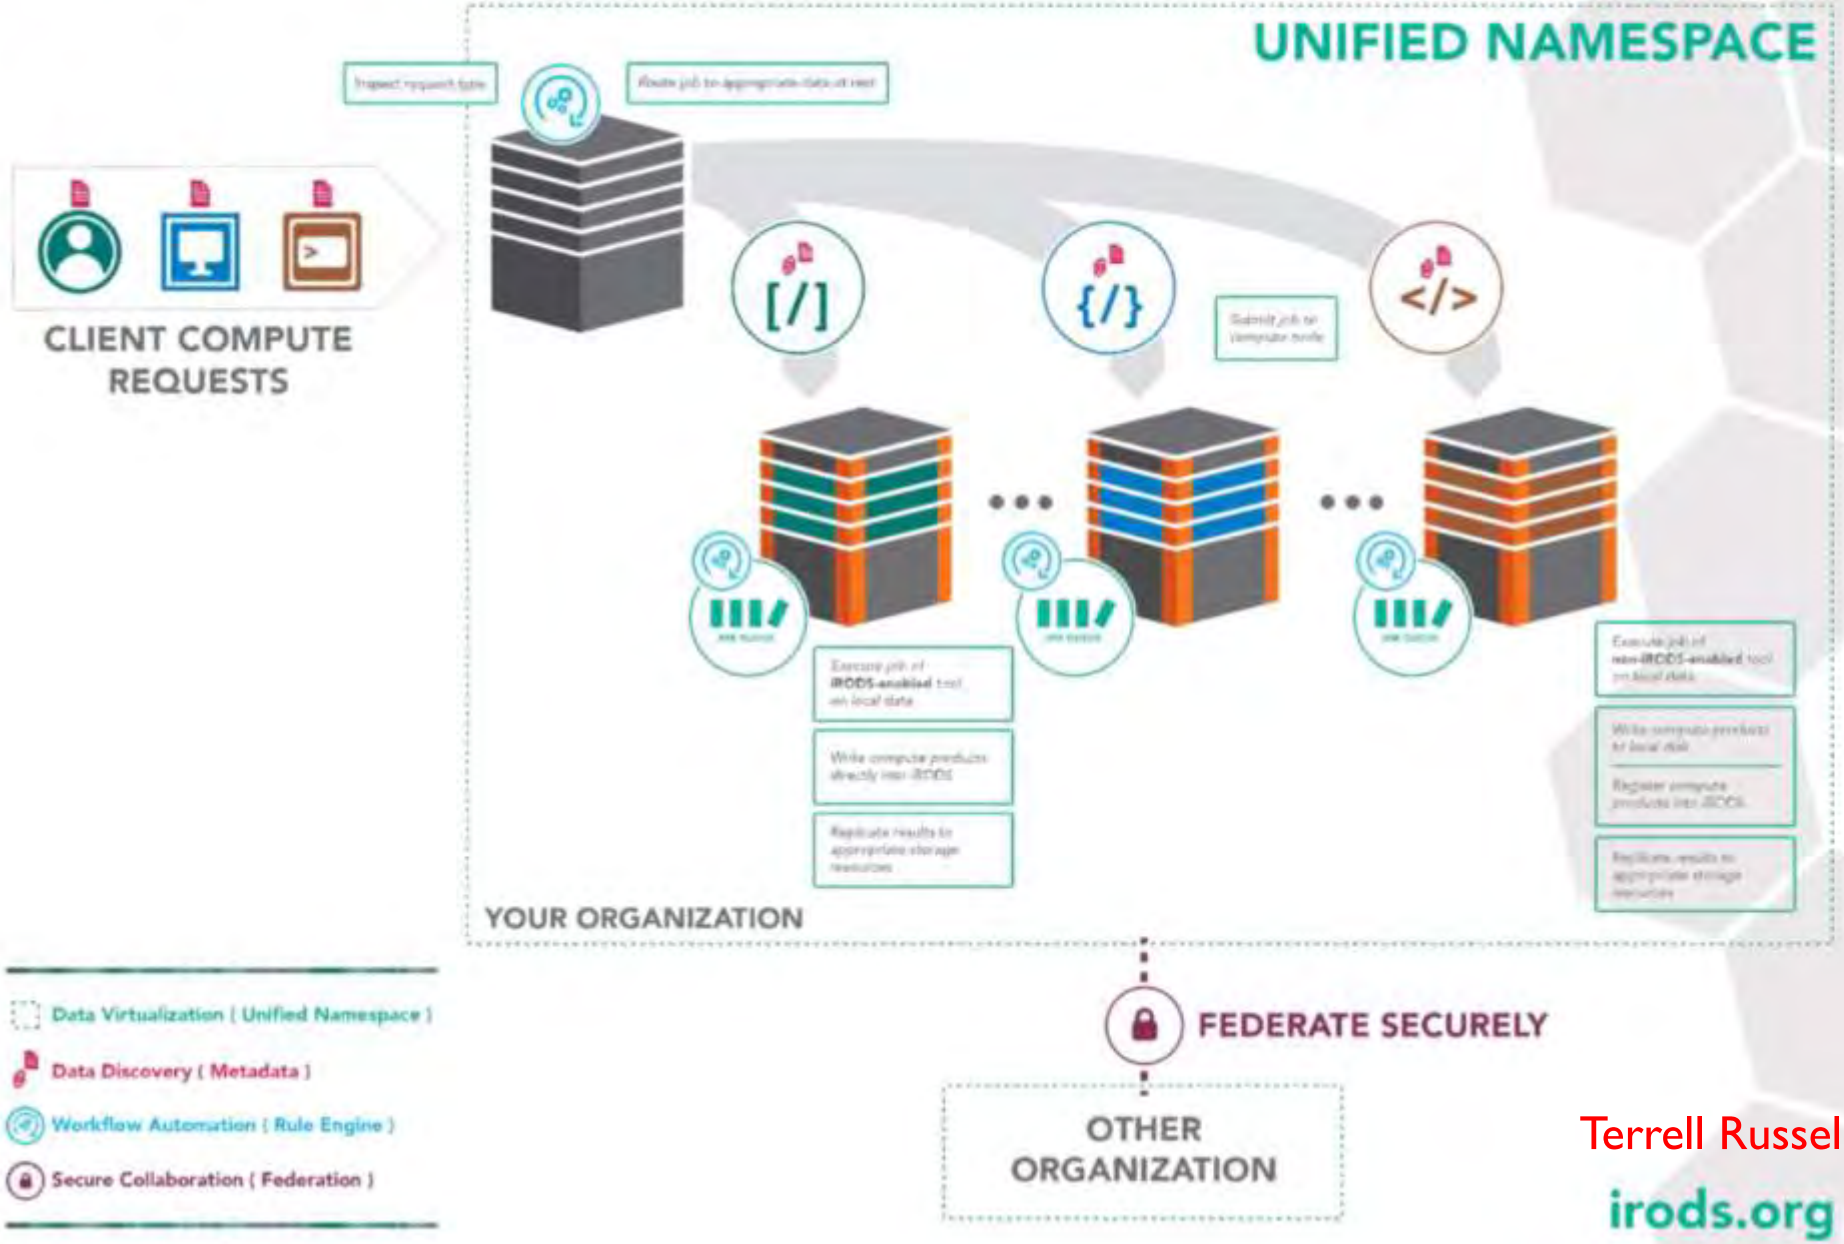

Terrell Russell

irods.org

Take Aways

iRODS

- Automatic, policy-based solutions are more future-proof as technology continues to change

- Having a programmatic interface (to the iRODS Rule Engine, via Policy Enforcement Points) means action(s) can be taken on your data based on the metadata:
  - Ingest
  - Metadata Extraction
  - Data Verification
  - Storage Tiering
  - Indexing
  - Publication
  - Auditing / Reporting

- Metadata templates allow for validation and verification
  - Match your domain-specific vocabulary and taxonomies
  - Reference outside standards
  - Prove compliance with required formats
  - Publish to make data discoverable

Terrell Russell