

DaRIS service description

DaRIS (Distributed and Reflective Informatics System) is a subject-oriented informatics framework built with the Mediaflux (<http://www.arcitecta.com>) data management platform. Its primary use is to supply biomedical imaging data management and integration with instruments and research computing infrastructure.

1. Overview

Collaboratively developed by the University of Melbourne, Arcitecta, Monash University and the University of Queensland, DaRIS has been successfully established in a number of operational deployments through the National Imaging Facility and other locations. DaRIS is built with Mediaflux, which has substantial use throughout the Australian research community as a generic data management capability. In particular, it is part of the RDSI-funded middleware stack provided at nodes around Australia (with unlimited license systems available). See the diagram in the Appendix for the architecture of the DaRIS/Mediaflux system.

The primary functionality feature set of DaRIS is:

- **Upload data**

DaRIS supports a variety of ways to upload data:

- Receiving DICOM metadata and images directly from scanners via the built in DICOM server¹
- Receiving raw data from scanners via DaRIS clients developed for Siemens MR, Siemens PET/CT and Bruker/Paravision.
- Importing other forms of data using the generic data upload capability in DaRIS web portal.

- **Organize your data**

DaRIS organizes data and metadata by a simple research oriented data model, where data are organized primarily by “Project” and “Subject”. The data model can be customized for specific research domains. In DaRIS

- All objects are managed together with the same provenance.
- Metadata are harvested automatically into indexed and searchable metadata during data ingesting. Other specialised meta-data can also be supplied for all objects in the data model.
- All interfaces are driven by the data model so that new functionality is provided for all users (no bespoke scripts and data structures)

- **View your data**

- Metadata and image can be viewed in DaRIS web portal. The web portal include simple image viewers for DICOM and NIFTI (2 options) images. Thumbnail images in browser supported formats can be uploaded for large generic images and displayed in the web portal.
- Other viewer can potentially be included
- It is also possible to integrate your viewer via the RESTful DaRIS API - this allows you to fetch data directly from DaRIS to your viewer

- **Download your data**

¹ The DICOM server has a rich set of configuration options, including de-identification of specified elements.

- Data can be downloaded to your local file system via the web portal.
- Data can be deposited to remote file storage via 'sinks', e.g. SSH servers, file systems, WebDAV servers.
- Downloads include optional on-the-fly format conversion.
- **Search your data**

All metadata in DaRIS are indexed and searchable. User can do metadata searches via the web portal. The search results can be exported to CSV/XML/PDF formats or downloaded as an archive.
- **Secure your data**
 - All transactions are audited, and data are checksummed to ensure data integrity.
 - Both data and metadata are versioned to track the history of change.
 - All the data (objects) are protected with ACLs. An ACL specifies which roles within the project are granted access to the data (object).
 - Project team members are grouped by roles. Each team member has a role within that project specifying their access.
 - Deleted data can be restored if soft-deletion is enabled.
 - Data can be replicated to other Mediaflux servers for disaster recovery.
 - Intrusion detection and authentication domain specific firewalls are supported by Mediaflux
- **Share your data**
 - Data can be transferred to remote (SSH, WebDAV) server sinks.
 - DICOM data can be sent to other DICOM servers by using the DICOM client in DaRIS.
 - Replication service can also be used to distribute data collections to other DaRIS instances.
 - OAI-PMH meta-data harvesting (e.g. ANDS Research Data Australia) can be enabled in DaRIS.
- **Process your data**
 - DaRIS can integrate with Nimrod/Kepler computation workflow to process data on the external computing resources.
 - Restful/Client APIs can be used on the compute nodes to download data from DaRIS and upload the processed data back to DaRIS.
- **Extend DaRIS**

Apart from the existing functionality, DaRIS is fully customizable and extensible. Based on the service-oriented framework of Mediaflux, you could

 - build your own plugin services (all DaRIS functionality is delivered by services)
 - build your own client applications with Java API
 - build your own clients with shell wrappers of the Mediaflux Java client.
 - DaRIS is now an open source project published at github.com. It is released under [BSD\(2-clause\) licence](#).

2. How the DaRIS service is delivered

2.1. Deployment options

VicNode offers flexible deployment options for DaRIS service. It can be deployed on your own Mediaflux server or we can host your project on the VicNode DaRIS (Mediaflux) server.

DaRIS software is open source and it is distributed under [BSD\(2-clause\) license](#) while its underlying platform, Mediaflux, is a proprietary software. VicNode also offers flexible Mediaflux licensing model to DaRIS customers.

The table below compares some of the requirements and features for the two deployment options:

| | Self-operated DaRIS | VicNode DaRIS |
|---|--|------------------------------|
| Purchase of Mediaflux license | Yes ² | No |
| Purchase of Mediaflux annual support | Yes ³ | Maybe from 2017 ⁴ |
| Infrastructure (server, storage and network, etc) provision required | Yes (NeCTAR cloud virtual machine is available) | No |
| Staff required to maintain the Mediaflux/DaRIS server | Yes | No |
| Optional use of RDSI storage | Yes | Yes |
| Replication to the disaster recovery server at VicNode | Yes | Yes |

2.2. VicNode DaRIS service

The VicNode DaRIS (Mediaflux) capability is available for all Victorian research institutions. It is hosted on a high performance computer (2 x 12-core Xeon, 192GB RAM, 4TB SSD, 2 x 10GbE), and it can be connected with the client's VicNode Market (RDSI) storage if required. It is supported jointly by Arcitecta (supplier of Mediaflux) and the VicNode data services team. At present, the service is funded through RDS/NCRIS to mid 2016. It is anticipated that the DaRIS service will be part of a wider imaging service (including Microscopy). Funding (regardless of method), will critically depend on uptake and demonstrated value add to research.

By hosting your projects on VicNode DaRIS, you eliminate the costs of Mediaflux licences, hardware infrastructure and support personnel and you have your project data managed by the high-performance data management system at VicNode. Additionally, your project can be replicated to the DR (Disaster Recovery) server at VicNode.

The DaRIS support at VicNode will help you identify the required information for your DaRIS projects:

- the data and metadata information (domain specific data model)
- the data replication requirements
- the network connectivity requirements
- how to integrate your instruments with DaRIS
- DICOM server settings (if required)
- other information that is required to define and deploy your project

² Mediaflux license packs are tiered; 5-user, 10-user, 25-user, 50-user, 100-user and so on.

³ Annual support (20% of license cost) can be activated/deactivated as desired. Mediaflux server upgrades are only available when annual support is activated

⁴ From 2017, this system will attract annual support. A business model to fund this has not yet been developed. It may include a modest contribution from stakeholders (ideally at the faculty level).

If at any time you choose to exit the service, your data and meta-data can be fully exported in its native form. If you do exit, all data would be deleted from VicNode's DaRIS system and the DR system (at a time jointly agreed to).

Storage for DaRIS Projects will be supplied as follows:

1. A starter storage of allocation of 100 GByte per project will be supplied
2. Once this quota is exhausted, users will need to apply for a storage allocation on VicNode. Please see <http://vicnode.org.au/> for more information on how to apply. At this point, the 100 GByte starter allocation is returned and all data moved to the user's new allocation.

2.3. Self-operated DaRIS service

If the VicNode DaRIS service is not applicable for you, you can also operate your own DaRIS instance. There are a number of flexible Mediaflux licensing options available as described below (licensing should also always be discussed directly with the supplier, Arcitecta - the following table is for information only).

2.3.1. Mediaflux licensing model

Firstly, you need to determine which type of Mediaflux licence suits you and purchase it from Arcitecta (or via VicNode). The following table describes available licence types:

| Mediaflux Licence Type | Description | Pricing | Type of Clients applies to |
|-------------------------------|---|---|--|
| Standard Mediaflux Licence | The standalone Mediaflux licence comes in tiers of 5, 10, 20, 50, 100 etc. user packs. | Contact Arcitecta for quotation. Educational discounts are available. Purchase is for 99 years with optional annual support (20% of license cost). | Groups or departments that need an in-house fully self-controlled solution. |
| DaRIS-Only Licence | The DaRIS-only licence comes in blocks of 5-users. Cheaper than standard licence, but allows Mediaflux server to run only DaRIS. The DaRIS support team (instead of Arcitecta) are the front line of support. The support is restricted to DaRIS related issues. | \$10,000 per 5 user blocks to purchase (99 years). 20% (\$2,000) per annum support. Can be upgraded to standard/full license at any point (just pay differential). | Groups wanting to run their own DaRIS instance. |
| Hub-Spoke Licence | The license model allows licences to be transferred from one system to another. As the VicNode Mediaflux has unlimited licences, it is possible to transfer licences (in the standard license packs 5,10,25...) from VicNode Mediaflux to another Mediaflux instance. However, this is only valid if VicNode continue paying for its unlimited licence. | The establishment license purchase cost is waived. Both the hub and the spoke must pay annual support (20% of license pack value). For example, if you transferred 10 licenses, then you pay 20% of the cost of 10 licenses per annum (which works out at about \$9000 for academic discount). | The institutional entity that operates the spoke must be the same entity that operates the hub. Very cost-effective way to leverage large central servers and retain autonomy. |

Mediaflux license packs are floating. If you have 20 users, 5 licenses may be sufficient depending upon access patterns. A license is consumed for a minimum of 30 minutes. The same user logging in multiple times from the same computer (IP) consumes only one license.

The DaRIS support team at VicNode can help you work out the licence type that suits you. See Section 4. for more details.

2.3.2. Deploy your own DaRIS

After the Mediaflux licence is purchased, you will need to allocate the computing resources, (e.g. server machine, storage, network) for Mediaflux and DaRIS.

The [NeCTAR research cloud](#) is available for the research institutions to deploy Mediaflux and DaRIS. The benefits of hosting DaRIS on a NeCTAR virtual machine include:

- high-speed network connection to AARNET and research computing facilities.
- integration with the large RDSI storage

The DaRIS support team at VicNode can help you configure the NeCTAR virtual machine and deploy Mediaflux and DaRIS on it.

If you choose to deploy Mediaflux and DaRIS on your own computing facility, the VicNode support team will also provide support for the installation and maintenance of Mediaflux and DaRIS.

Regular consultation sessions to help you get the most for your researchers are a part of the VicNode DaRIS service.

2.3.3. Setup your DaRIS projects

After Mediaflux and DaRIS are deployed, the DaRIS support team can help initialize your projects by defining the domain specific metadata and customizing DaRIS clients. They will help you build your own domain-specific package by working from a generic template package.

2.4. DaRIS helpdesk

The DaRIS helpdesk at VicNode is available for all DaRIS customers to provide support for DaRIS end users, administrators and developers. Currently, DaRIS support can be contacted by email: daris-support@lists.unimelb.edu.au.

For groups who operate their own DaRIS servers, the daris-admin google discussion group, <https://groups.google.com/forum/#!forum/daris-admin> is available for DaRIS administrators.

For developers, DaRIS source code is available at <https://github.com/uom-daris>. Furthermore, VicNode DaRIS support team can work with DaRIS customers to:

- create their own domain specific plugin package, which includes the definition of the domain specific metadata and customization of the services.
- develop or customize client applications to handle domain specific data formats (as requested).

2.5. Try DaRIS

You are able to trial the DaRIS service at VicNode at no cost. A test project and trial user accounts will be created and you will be supported to utilise them effectively.

To apply for DaRIS trial, please contact daris-support@lists.unimelb.edu.au

3. Apply for DaRIS service

The DaRIS team at VicNode can help you apply for the VicNode DaRIS service or to deploy your own DaRIS, including the applications for Mediaflux licence, NeCTAR cloud and RDSI storage.

To apply for DaRIS service, please contact daris-support@lists.unimelb.edu.au

4. To find out more

- **DaRIS project:**
 - Documentation: <http://nsp.nectar.org.au/wiki-its-r/doku.php?id=start>
 - Source code: <https://github.com/uom-daris>
 - Google group: <https://groups.google.com/forum/#!forum/daris-admin>
- **RDSI storage** at VicNode: <http://vicnode.org.au/>
- **Arcitecta/Mediaflux**: <http://www.arcitecta.com/>
- **NeCTAR research cloud**: <https://www.nectar.org.au/>

| Version | Date | Who |
|----------------|-------------|-----------------------------|
| 1.0 | 07 Aug 2015 | Wilson Liu and Neil Killeen |

Appendix - DaRIS/Mediaflux Architecture

