

# SuperDARN Data Distribution Working Group Charter

Kevin Krieger  
Kevin Sterne  
Ethan Miller

November 2013

## Contents

<b>1 Purpose</b>	<b>2</b>
<b>2 Holding Institution</b>	<b>2</b>
<b>3 Distribution Institution</b>	<b>2</b>
<b>4 Membership</b>	<b>2</b>
<b>5 Communications</b>	<b>3</b>
<b>6 Responsibilities</b>	<b>3</b>
<b>7 Immediate Action Items</b>	<b>4</b>
<b>8 Points for Consideration</b>	<b>4</b>
<b>A Brief account of the functions of the holding institution for the data distribution that were previously performed at JHU/APL</b>	<b>5</b>
A.1 Notes . . . . .	5
<b>B An account of the functions of the new holding institution at Virginia Tech</b>	<b>7</b>
B.1 Tracking Issues with Coverage and Data Quality . . . . .	8
<b>C An account of the system in place at the data distribution institution - University of Saskatchewan</b>	<b>9</b>
C.1 Previous system overview . . . . .	9
C.2 Current system overview . . . . .	9
C.3 Mirror overview . . . . .	10
C.4 User accounts . . . . .	11
C.5 SFTP details . . . . .	11
C.6 Daily operations . . . . .	12

# 1 Purpose

This working group (WG) is charged with maintaining the protocols for transferring RAWACF and DAT files from the PI institutions, first to a holding institution (Virginia Tech) and ultimately to a distribution institution (University of Saskatchewan) where they are merged and copied onto hard drives for distribution. The cadence for distribution of the drives is every four months (Jan-Apr, May-August, September-December). This working group was created after the SuperDARN workshop in Moose Jaw, Saskatchewan in 2013. The group is also charged with maintaining an augmentation to the hard drive distribution, any networked SuperDARN Data Mirrors.

# 2 Holding Institution

The holding institution prior to October 1st, 2013 was JHU/APL. See Appendix A for a description of the system that was in place at JHU/APL.

The holding institution beginning October 1st, 2013 is Virginia Tech. See Appendix B for a description of the system in place at Virginia Tech.

# 3 Distribution Institution

The distribution institution is the University of Saskatchewan. See Appendix C for a description of the current system in place at the University of Saskatchewan, and a brief description of the previous system.

# 4 Membership

The membership of the working group should consist of several types of members:

1. Those responsible for providing data to the holding institution.
2. Those responsible for accessing data from data mirrors to download to their PI institution.
3. Those responsible for receiving hard drives from the four month distribution.
4. Others who wish to observe.

The current membership includes:

- Chair: Kevin Krieger University of Saskatchewan
- Co-Chair: Kevin Sterne Virginia Tech
- Simon Shepherd Dartmouth
- Akira Sessai Yukimatu NIPR
- Kehe Wang IPS
- Andrew McDonald LaTrobe
- Paul Breen BAS
- Bill Bristow UAF
- Jef Spaleta UAF
- Jianjun Liu PRIC
- Judy Stephenson Durban

- Tomoaki Hori STELAB
- Nozomu Nishitani STELAB
- Aurlie Marchaudon IRAP/CNRS
- Nigel Wade Leicester
- Tsutomu Nagatsuma NICT
- Keisuke Hosokawa DICE
- Ethan Miller JHU/APL
- Ermanno Amata IAPS-INAF
- Tom Kane LaTrobe
- Hong-Qiao Hu PRIC
- Xavier Vallieres CNRS
- Chris Thomas Leicester
- Manabu Kunitake NICT

Observers:

- Mike Ruohoniemi Virginia Tech
- Kathryn McWilliams University of Saskatchewan
- Dieter Andre University of Saskatchewan
- Jean-Pierre St.Maurice University of Saskatchewan

## 5 Communications

Communication shall primarily be done over email via the Data-Distribution Working-Group (Darn-ddwg) mailing list set up by Julian Thornhill, maintained by Leicester University. To send email to the list, send mail to [darn-ddwg@ion.le.ac.uk](mailto:darn-ddwg@ion.le.ac.uk).

A person wishing to subscribe to the Data Distribution emailing list should contact the Chair (or co-Chairs) or the current PI institutions member about subscribing to this list. Subscribing to the list should be only at the approval of the Chair (or co-Chairs) so that they are aware of an addition or change within a PI institution.

## 6 Responsibilities

- Chair/Co-Chair: Provide updates on the status of the SuperDARN data distribution system to the members of the ddwg. Operate and maintain the SuperDARN holding and Distribution institution systems. Track down gaps in data and resolve them, either by providing an explanation for why the data is missing, or retrieving it from the PI institution responsible if it is not a real gap.
- Members responsible for providing data to the holding institution: Provide data in a timely and preferably automated manner to the holding institution, either by 'pushing' the data or by placing it in a holding area for the holding institution to 'pull'. Help with gap resolution when requested.
- Members responsible for accessing data from mirrors: Set up and maintain a way of downloading data from the mirrors in a timely fashion for their PI institution. Provide feedback on the mirror setup.

- Members responsible for receiving distributed hard drives: Make sure the hard drive received contains the data that it is labeled for, and is readable.
- Observers: Provide feedback or input on discussions.

## 7 Immediate Action Items

1. An introduction email shall be sent out by the Chair/Co-Chair to the Darn-ddwg mailing list, with this charter.
2. An email shall be sent out detailing the current status of data distribution, including the mirror.
3. Documentation shall be done on the Primary SuperDARN Data Mirror.
4. Discussion among all members should happen regarding the Points for Consideration below.

## 8 Points for Consideration

1. Are there problems with accessing files from certain institutions?
2. How do we become aware of, and deal with, gaps in the distribution?
3. Have bad files been passed to the general distribution? Should these be identified and corrected, deleted, or somehow flagged? (Should be less of a problem after Oct. 1, 2013)
4. How do we make rawacf files available on a next-day (or better) basis?
5. Do we have issues with the amount of data and storage capacity?
6. The files currently distributed by U SAs are apparently not identical to those that were provided by the PI institutions. Why is this, and is it a problem?
7. Should the Data Mirror set up at U SAs eventually replace distribution of data by hard drives?
8. How do we ensure that Virginia Tech has the most complete data archive? There are known gaps within the Virginia Tech data archive across several radars (April 29 May 6, 2008; July 30, 2008; December 30, 2008)
9. How do we ensure that much older data that is corrected by an institution is distributed throughout the SuperDARN community? For example, if a flaw is found in 2009 Blackstone data and it is corrected, how do we ensure the corrected data is distributed throughout the PI institutions?
10. Building upon #4, how can the community quickly scrutinize/analyze data to ensure, at least at first glance, proper radar operation? If a days worth of data contains 10 return-cells per beam per hour, is the radar broken? Hopefully this will help keep faulty data from quickly being distributed and lessen the need for my #2.
11. How do we process different/irregular file formats into the distribution? The University of Alaska Fairbanks radars have begun producing multi-channel data that cannot utilize the channel variable in the dmap header due to stereo-ops logic in the rst code that breaks. So, to keep things sensible, the channel information is being stored into the filename. As well, older radars still produce data in the old fit and dat file formats. Which do we use as a basis for converting to the new file format?

## A Brief account of the functions of the holding institution for the data distribution that were previously performed at JHU/APL

1. The data distribution tasks are handled day-to-day by Stu Nylund.
2. Rawacf and fitacf files (and wals iqdat files) are generated at the remote radar stations and collected, either daily via the network or periodically via disk delivery, at remote sites such as VT, Fairbanks, CNRS, etc.
3. At 0200 LT each night, APLs morpheus runs a cron job to download all the latest rawacf and fitacf files and processes them (creating a log file in the process). The stages of this processing are as follows
  - (a) Download: Each data file is downloaded from the remote site to a station-specific subdirectory under the incoming directory and, as record, an zero-lengthed file ([filename].cat) is created in a station-specific subdirectory under the cat directory. For some remote sites, the remote files are deleted once they are downloaded to APL.
  - (b) Backup: All downloaded files are copied to a station-specific subdirectory under the backup directory. This unchanged files serves as a record of the stations version of the original file.
  - (c) Tests: Most remote sites data files are posted as compress files (gzip or bzip2), so a next step is to check that the compress files is good (i.e., as a test of a valid network transfer). If the compressed file is bad, the file is moved to a station-specific subdirectory under the bad directory. Similarly, the file is moved to the station-specific bad directory if it is empty, had a problem in naming or in being stored locally (i.e., an APL disk becomes full). Error entries for each of these cases are made in a station-specific log file.
  - (d) Conversion: The file is next modified/converted slightly (e.g., headers properly filled in, etc) I am not familiar with this conversion step, and Rob had talked about changing/eliminating it. If any problem is detected, the file is put in a station-specific subdirectory under the bad directory.
  - (e) Storage: The converted file is next copied to a type-specific subdirectory (rawacf, fitacf, iqdat) under the data directory for storage.
  - (f) Upload: The uploading of files to Saskatchewan is done in two steps an automatic preliminary step under the process cron job and a manual step performed periodically
    - i. Automatic: The converted file is next copied to a type-specific subdirectory (new.rawacf, new.fitacf) under the outgoing directory for accumulation and subsequent uploading to Saskatchewan.
    - ii. Manual: An upload is later manually initiated by scripts (via upload.rawacf and upload.fitacf); under either script, the files are moved to a type-specific subdirectory (rawacf, fitacf) under the outgoing directory, all files in this subdirectory are rsync uploaded to Saskatchewan, and finally each file is then moved to a type-specific subdirectory (arc.rawacf, arc.fitacf) under the outgoing directory.

### A.1 Notes

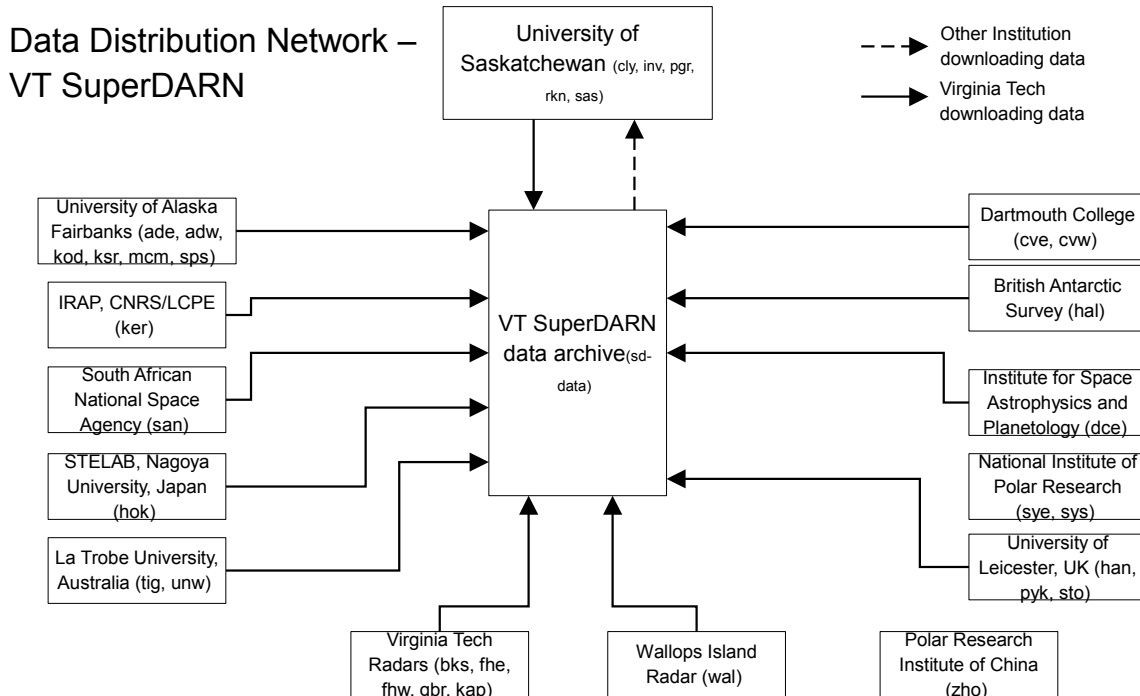
- As far as possible RAWACF files should be transferred daily to the holding institution
- The files are compressed before transfer to the holding institution (bzip2, gzip, others)
- The standard for data transfer is a script that runs from the holding institution that pulls RAWACF files from a directory designated by the PI via sftp or scp; if a ssh-based data transfer protocol is not possible, ftp could be used. Alternatively, the PI might opt to push his files into a directory designated by the holding institution for this purpose
- U Sas may upload files from the holding institution at any time
- Checking of files for completeness before merging is carried out at U Sas; the data distribution working group member of each PI institution is asked to confirm or correct apparent gaps

- Beginning Oct. 1, 2013 the holding institution (Virginia Tech) will perform additional checks on the files for integrity as they arrive and notify PIs of any issues
- If a regular, automated connection to the holding institution cannot be made from a PI institution, measures will be taken to route the data transfers to the holding institution through a relaying institution.

## B An account of the functions of the new holding institution at Virginia Tech

With the extension of the mid-latitude chain of radars across North America, the Virginia Tech SuperDARN group has pursued collecting data from these radars as well as other radars on a daily basis. These daily downloads have been set up to mostly run in an automated fashion so that it does not tax a person to run the scripts.

So far, the downloads are setup such that our data archive accesses another institutions archive or their staging area which may contain several radars data. Several of these data download connections are illustrated in the [graphic](#) below.



Here it should be noted that most of these data transfers are files in the rawacf file format. The exceptions to this include dat files that are being downloaded from La Trobe University for the Unwin radar (unw) and processed into rawacf files locally at Virginia Tech. Also, fitacf files are being downloaded, but not archived, from the University of Alaska Fairbanks for the Kodiak (kod), McMurdo (mcm), and South Pole Station (sps) radars. Also of note here, the downloads from IRAP for the Kerguelen (ker) are only 8 hours of rawacf files on a daily basis. Bandwidth limitations restrict a full days worth of downloads to occur. However, the full data set for the Kerguelen radar is shipped out every few months and the same download procedure will be used to bring the full data set to Virginia Tech.

As far as radars that Virginia Tech manages (bks, fhe, fhw, gbr, kap), data is downloaded from these radars on a bi-hourly basis and staged up for the University of Saskatchewan group which in turn downloads the files on a daily basis. While Virginia Tech downloads both rawacf and iqdat for most of these radars, only rawacf files are staged up for the University of Saskatchewan group to download.

All of the daily downloads are easily tracked by using our data inventory webpage at:

<http://vt.superdarn.org/tiki-index.php?page=Data+Inventory>

We show the last 7 days worth of data as sometimes links to radars go down and a day or two of lag can occur in acquiring the data. A more complete view of the data from all of the radars in the SuperDARN network can be viewed in the graphics further down in the page. In list form, these are the radars that Virginia Tech is currently downloading:

Institution Downloading From	Radars
Directly from radars	bks, fhe, fhw, gbr, kap, wal
University of Alaska Fairbanks	ade, adw, kod, ksr, mcm, sps
Dartmouth College	cve, cvw
STELab / Nagoya University	hok
University of Leicester	han, pyk, sto
La Trobe University	tig, unw
British Antarctic Survey	hal
IRAP, CNRS/LCPE	ker
Institute for Space Astrophysics and Planetology	dce
Durban	san
University of Saskatchewan	cly, inv, pgr, rkn, sas

## B.1 Tracking Issues with Coverage and Data Quality

Even with all of these downloads and the tools to view and verify the data, several problems still exist in distributing the data as well as known problems with data. These include:

- No Zhongshan (zho) rawacf files since late June 2010, no Zhongshan (zho) fitacf files since early February 2011.
- Channel information recording convention for Adak East, Adak West, McMurdo, Kodiak, and South Pole Station. Here the channel information is stored in the file name, instead of within the data file as currently used by University of Leicester radar systems. The rawacf files distributed by JHU/APL do not appear to have properly combined data showing multi-channel data.
- Bad velocities for King Salmon data between October 14 to October 25 16:17 utc, 2011
- First range errors noted at some University of Saskatchewan radars as well as Blackstone due to receiver driver sampling errors. Steps to correct and detect this issue have been taken by the Virginia Tech and University of Saskatchewan through varying methods.
- Unsure of the completeness of our data archive. With some groups that allow access to the complete dataset for a particular radar, some effort has been put forth to check that the data present in our archive matches the other institutions archive. For now, this has been crudely done by matching file names, if a file name exists on another institutions archive then we download it to our archive.



## C An account of the system in place at the data distribution institution - University of Saskatchewan

### C.1 Previous system overview

Prior to October 1st, 2013, JHU/APL (Robin Barnes/Stu Nylund) downloaded newly posted data from approximately 16 institutions nightly. This data was then processed through their scripts before they were eventually posted to the University of Saskatchewan's sdcopy server. Once four months worth of data (January to April, May to August, or September to December) are uploaded to sdcopy by JHU/APL, the data manager would write 17 hard-drives with copies of the four months worth of data. The hard drives were then shipped out to 16 institutions and one would be kept at the University of Saskatchewan for an extra backup.

In preparation for a distribution of data, the data manager at the UofS prepares 17 external hard drives by formatting them with an ext3 filesystem, partitioning them, and creating several directories for rawacf and fitacf data files to go. Fitacf files would be placed on the hard drive for radars that could not reliably transfer rawacf files.

After four months worth of data has been transferred to sdcopy by JHU/APL, the data is copied to six hard drives at a time (presently using USB external hard drives). This takes about two work days. After the data is copied to the six connected hard drives, the data on the hard drives is read and checked for consistency with the original data on sdcopy. This diffing takes roughly another work day.

Alongside the preparation of the hard drives, gap resolution would be done. The data manager would run a script on all the files, for all the radars over the past four month period. The output of the script is in the format:


```
20120318.14.rkn -- 20120318.22.rkn
20120319.00.rkn -- 20120319.14.rkn
20120321.16.rkn --
```

This format means that there are no rawacf data files during the time period March 18th, at 14:00 UTC to March 19th at 14:00 UTC. Then there are data files for March 19th 16:00 UTC until March 21st at 14:00 UTC. And there are no data files for March 21st at 16:00 UTC.

Emails are sent out to the individuals responsible for each radar, and they respond with information regarding the gaps, if they are real or not. If the gaps are not real, then they provide the data files to be included in the next distribution of hard drives.

### C.2 Current system overview

The current system in place at the University of Saskatchewan is transitioning to using Virginia Tech for data. Hard drives are still planned for distribution every 4 months. Refer to appendix B for a description of the Virginia Tech system. Data is downloaded daily from VT into a holding area on our superdarn-cssdp server. Then, a gatekeeper script goes through the holding area looking for new data. Any new data is added to the appropriate place on the mirror. If the gatekeeper sees any errors (such as reposted data, or data that has the same file name but different checksum as that on the mirror) then it does not add it to the mirror. Upon finding an error, the gatekeeper informs the data manager through email, as human intervention is required at that point.



### C.3 Mirror overview

The SuperDARN Primary Data Mirror is implemented using WestGrid storage space. WestGrid is part of a consortium of High Performance Computing (HPC) resources available to academic researchers. The WestGrid storage space is accessed on a desktop computer running Linux in the SuperDARN computer room, and is mounted as a Networked File System (NFS). In order to ensure adequate security, the protocol used is the Secure File Transfer Protocol (sftp), and it runs in a segregated environment on the Linux computer.

The segregation allows the users of the SuperDARN Primary Data Mirror read-only access to the data, and only the directories that contain the data. In this way, the rest of the file system is inaccessible to users.

See [this illustration](#) for a graphical view of the directory structure of the SuperDARN data mirror.

A short description of the directory structure:

- `/sddata` This is the root of the mirror.
- `/sddata/master.hashes` This file is updated daily and contains a line for every `YYYYMM.hashes` file on the mirror. It is useful for determining what months have new or changed data, as the hash for that specific hashes file will change. An example usage of this is to keep a local copy of this master.hashes file, and periodically compare your local copy with the server's copy to see what months you need to synchronize with.
- `/sddata/raw` This is the directory for 'raw' data (currently, only 'rawacf' files).
- `/sddata/raw/2010` Directory for 'raw' 2010 data.
- `/sddata/raw/2010/04` Directory for 'raw' April 2010 data.
- `/sddata/raw/2010/04/201004.hashes` A list of all the hashes for files in the `/sddata/raw/2010/04/` directory. It is up-to-date and is used by the example synchronization script sent out previously, as well as the gatekeeper to the mirror. Every time a new or updated file is placed in the `/sddata/raw/2010/04/` directory, it will be updated. It is not sorted.

## C.4 User accounts

In order to set up a user account, contact Kevin Krieger at the University of Saskatchewan. Three things are required in order to set up a user account:

- A static IP address OR a hostname of the machine you would like to connect to the sftp server from.
- A username you would like.
- A Public key from a public/private key pair.

In a Windows environment, Public-Private key pairs can be generated using the program `puttygen.exe`. After generating the key pair, if you hit the button "Save Public Key", it doesn't save it in the correct format. Comments need to be removed and newlines need to be removed, and the text `ssh-rsa` and a space need to be placed in the beginning of the line. Instead, copy the public key from the text box and save it to a file.

In a Linux environment, the program `ssh-keygen` will by default generate a 2048 bit RSA key pair when run from the command line without arguments. It will ask several questions and generate two files, which are in the appropriate format. A passphrase is not required and generally left blank.

If you are running a Windows environment, then a program such as WinSCP can be used. If you are running in a Linux environment, then there is the command line program 'sftp' that I believe is generally packaged with linux distributions. Otherwise, the program FileZilla also supports sftp.

## C.5 SFTP details

- The login shells for the mirror users will be set to `/bin/false` (so there's no way to get to a shell)
- Maximum concurrent sessions per user will be 2 initially, but if this proves to be an issue, we can definitely change things.
- The maximum attempted logins (OpenSSH option `MaxStartups`) is set to a fairly low number, I don't see this being an issue, but again it can be revisited.
- `LoginGraceTime` is set to 1 minute, this means that if a connection to the server is sitting and unauthenticated, this timer will start and will end the unauthenticated session after the time limit is reached.

The permissions and ownership of the files in the chroot directory affect the commands that are available to a user. The following commands are available:

- `df` - This shows disk usage
- `ls` - This shows a listing of the files, just like a standard `ls`
- `cd` - Change directory
- `progress` - Toggle a progress meter on and off
- `pwd` - Print the working directory
- `quit` - Quit the sftp session
- `exit` - Quit the sftp session
- `bye` - Quit the sftp session

The following commands are not available because of permissions that the users do not have (namely write permissions or execute permissions).

- `chgrp`
- `chmod`
- `chown`
- `mkdir`
- `ln`
- `put`
- `rm`
- `rmdir`
- `rename`
- `symlink`

## C.6 Daily operations

There are several scripts that run on a daily basis on the SuperDARN mirror computer.

- `download_vt_data /path/to/holding/area RADARID`

This script is run once each day for each RADARID (3 letter radar code, such as 'ade' or 'cly'). The script will check for any new data available on the VT server, and download it. It will also check that the download completed successfully, and removes the VT holding area file if it did. If there are any issues, an email is generated to get some human intervention.

- `gatekeeper /path/to/holding/area RADARID /path/to/mirror/root/`

This script is run daily for each RADARID (3 letter radar code, such as 'ade' or 'cly'). The script will go through each file in the holding area for the specific radar and do the following:

1. Check the file for decompression errors using `bunzip2 -t`
2. Check if the file already exists on the mirror
3. If the file exists on the mirror already and is different, send email and don't move file.

4. If the file exists on the mirror already but is the same, just remove the holding area file and send email.
  5. If the file doesn't exist already on the mirror, first move the file to the mirror, then create a hash of the file and add it to the existing hashes file for the appropriate year and month.
- **generateMasterHashes**

This script will run once a day and update the *master.hashes* file under the */sddata/* mirror directory.