



# INTEGRATED DISSEMINATION PROGRAM – TECHNICAL APPLICATION STANDARDS

Version 1.3

This document explains the technical requirements for applications running on the IDP.  
The Development Organization must conform its applications to these standards.

## TABLE OF CONTENTS

<b>Revision History</b>	<b>2</b>
<b>Hardware Design Standards</b>	<b>3</b>
<b>Software Design Standards</b>	<b>4</b>
<b>Software Syntax Standards</b>	<b>8</b>
<b>Documentation Standards</b>	<b>9</b>
<b>Testing and Validation Standards</b>	<b>11</b>
<b>DataFlow</b>	<b>12</b>
<b>Configuration Management</b>	<b>13</b>
<b>Appendix I - Summary</b>	<b>15</b>

## Revision History

Version	Date	Author(s)	Revision
0.5	9/2018	A. Al-Mallah	Created
1.0	7/30/2019	C. Klemmer & S Earle	Document Released
1.1	9/6/2019	OBT	Minor updates, including logging and environments
1.1	1/8/2020	OBT	Included version specific software, embedding code practices, and handling site differences
1.1	1/8/2020	C Klemmer	Released
1.2	8/5/2020	C Klemmer	Updated to include database testing, user sign on, languages, update WWA failover, load testing expansion
1.3	3/8/2021	J Huber	Updated Monitoring section, added general coding standards to software design section, Added end-to-end testing under Testing and Validation, clarified supported versions of databases. Added LDM verbosity requirement

This document describes the technical requirements for the Development Organization (DevOrg) to follow when designing, upgrading and optimizing the applications that run on IDP. Prior to any code handoff to NCO this document should be reviewed and all necessary changes should be applied. If for any reason you believe your application does not comply with anything within this document please speak to the On Boarding Team (OBT) Lead. They will work with you to determine a path forward into operations.

## Hardware Design Standards

### Cluster Configurations

We will support 2 clusters, one will be dedicated for applications that directly support dissemination of Watches, Warnings, and Advisories (WWA) and the second for everything else unrelated to WWAs.

### Hardware Specifications

At a minimum, the development organization will have a hardware specification for each Virtual Machine (VM) that contains:

- Required number of Virtual Central Processing Units (CPU)s
- Required Virtual Memory
- Required RAM disk
- I/O traffic footprint

Both CPU and memory requirements should be built with an expected percentage utilized at 80% under maximum load.

A note on the number of CPUs - Recommended 2 to 4 CPUs, 1 core per CPU. There are 24 CPUs per blade, and the hypervisor is constantly swapping different VMs in and out so that each of them gets some CPU cycles. Our experiments suggest that once one exceeds 4 CPUs the amount of extra processing power gained becomes questionable, because once the VM is not active on the blade, it becomes more and

more difficult to put it back in, since several other VMs must be taken off at the same time to make room for it. The exceptions are VMs that run databases -- they can have 6 CPUs.

## Operating Systems

The following operational operating system is supported on the IDP infrastructure:

- Red Hat Enterprise Linux: Version 8

If you believe that your application requires an operating system other than Linux you must obtain an exception by the NCO Director.

## Disk Specifications

The following specifications must be provided for each disk an application accesses:

- Requirement for any local disk space
- For shared NFS mounts and local mounts the following information must be specified in the software requirements:
  - X Name that will reside under:  
*/common/data/apps/application\_name/tier/X*  
Where *application\_name* is how you are identified on IDP  
Where *tier* is either “dev”, “qa”, or “ops”
  - Size
  - Maximum number of files (inodes)
  - Latency requirements
  - Peak I/O generated by application

## Software Design Standards

### System Components

Do not embed any code that should be in an application centralized repository. Doing so prevents security patching from being executed. This includes any software that is downloaded through the official RedHat Network Channel. Note that when it comes to designing any recurring process a check shall be made to ensure that an instance of itself is not already running before it kicks off.

### Site differences

It is required to keep system and application configurations the same between College Park and Boulder sites. If they need to be different (eg:due to IPs or VM names, etc), these differences have to be documented in the top of the configuration file and easy to find. Do not embed differences within the configurations. Where possible, keep the

scripts identical between sites and use a reference pointer to a centralized configuration file.

## General coding

This section applies to all scripts and compiled code used by the application

- No code will have hard coded paths. All paths to get to input or output data will be read from environment variables or configuration files
- Variables used to define an environment (ops|qa|dev) will be read from config files
- Any site specific definitions also will be read from config files
- Application and scripts must be able to complete processing if an error condition is encountered.

## Scripting

This section applies to those scripts used by the application or used to maintain the environment

- The first line of the script will always identify the SHELL being used
- Every script must contain a clearly demarcated comment section that explains what the script does, what its expected inputs are, what the outputs are and what the errors/exceptions thrown are.
- All scripts (bash, perl, python) permissions will be set as 775
- Any script utilizing command line utilities as-is must use BASH
- Scripts must always return a predefined status value
- Errors/Exceptions are echoed to stderr and written to log files
- Any function that is not both obvious and short must be commented
- Use blank lines between blocks to improve readability. Indentation is two spaces, don't use tabs.

## Languages

The application shall be implemented in any of the supported tools/languages/frameworks:

Compiled Code	C++, Java
Web Application	Java, PHP, Perl, JavaScript, HTML5
Interpreted Code	Bash, Perl, Python (Anaconda)

- If any application were to use an enhanced framework in addition to one of the above listed languages, an appropriate justification needs to be specified. Support agreements for such frameworks need to be in place before they can be used in IDP and discussed with the OBT lead.
- These external frameworks may need further review by the NCO security team and the OBT. The most likely questions regarding external frameworks are their support structure, and security aspects. Another consideration would be to evaluate open source equivalents, if available

## Shared Software

If applications are leveraging shared software between them it needs to be kept in a common NFS area where both systems can point to it. It is not to be installed locally. A minimum of 2 versions of the software will be kept at all times. Applications will migrate to the latest version with a planned upgrade.

## Web Services

The following is a list of the web servers that are supported on the IDP infrastructure:

- Apache: Version 2.4
- TomCat: Version 7

IDP servers data via HTTPS or FTP front ends. All IDP web applications reside in a shared front end infrastructure of load balanced web servers. If you believe your application requires dedicated front end infrastructure, appropriate justification needs to be provided to the OBT lead. All websites will be hosted out of Akamai unless noted by OBT.

## Queuing

Use standardized mechanisms for common applications such as queueing and clustering, as opposed to using custom applications to achieve the same result. The following queuing mechanisms are allowed on the IDP infrastructure:

- Local Data Manager (LDM) Version 6.13
- RabbitMQ (<https://www.rabbitmq.com/>)

## Databases

All databases need to be able to drop all tables and start from scratch at any time, without any negative impacts to the functionality or stability of the application.

The following database servers are supported:

## Relational

- Postgres Version 9.0 or higher
- Mysql Version 5.6.x

## Non-Relational

- MongoDB Version 3.4

For new system builds, the below versions are supported:

## Relational

- Postgres Version 12.x or higher
- Mysql Version 5.7 or higher6

## Non-Relational

- MongoDB Version 4.4 or higher

## Compilers/IDEs

The following compilers are allowed on the IDP infrastructure:

- GNU open source suite (g++) Version 4.8.5
- Java (javac) OpenJDK Version 1.8.0\_x
- Eclipse (IDE)

## Build Tools

The following build tools are allowed on IDP applications:

- Makefile
- Jenkins
- Maven

## Third Party Software

Before an application is onboarded:

- Developers must identify any 3rd party software used. Which version and the vendor support associated with it.
- Any 3rd party software used has to be approved by the security team
- Include any licensing costs which would have to be approved ahead of time by the NCO director. License funding will be provided by the devorg, or agreed to be covered by DIS.
- Open source tools must be considered, if available and a support agreement with a vendor established.
- Appropriate justification for the tool selected for a purpose must be provided. This could be a business case for using a specific software package



## Unsupported(blacklist) software

Software that is not allowed:

- Stand Alone Anaconda
- Oracle Java

## Failovers

Applications must be designed to be able to failover to the non primary site within 15 minutes with no data loss and minimal data latency. An application that distributes Watches Warnings and Advisories (WWAs) must be able to failover in under 5 minutes with no data loss.

The failover must be designed to require minimal human intervention. The use of an automated “one button” script which switches the primary DNS to fail-over successfully is required. The application must be able to detect and report on whether the failover has worked properly within 10 minutes of the failover. An end to end test to confirm the products are being produced and received by the intended application(s) must be included.

## Logging

Standardized application loggers are to be used and follow these guidelines:

- Logs will be kept for a 7 day retention
  - Log retention periods needs to be configurable in a script
- Every independent piece of the application will have its own log file
- All logs, except current day’s log, must be compressed using gzip
- The level of logging must be able to be changed in real time without recompiling the application. Instructions are to be provided on how to change the log verbosity.
- Logs will be written to the dedicated NetApp or local application volume; never under the local root volume.
- All application logs will have read/write permissions by the application user account and read permissions by group/world.
- Developers must use log rotate coding for a maximum of daily rotation
- All critical transactions will be included in the log output
  - Any “deletes,” “writes” or “moves”
  - Any application that transmits data offsite must include a timestep of when the file was successfully sent
- Any application that transmits WWA’s must have a more verbose output to fully

describe the reason for the logged transaction.

- Any application that cannot use the system utility 'log rotate' must include a configuration file that determines where logs are stored, how often they rotate, and how long they are saved.
- LDM logs must be set to verbose by utilizing the -v option at start up.

### Format of Log Cleanup

- BASH script will reside under \$HOME/config called "log\_cleanup"  

```
>retention_period=7 (defined as number of days)
>find $HOME/logs -mtime -${retention_period} -type f -exec rm {} \;
```

### Format of Log Contents

- Log filename must follow this pattern:  
app\_function\_YYYYMMDD.log
- Include date and time to the second for every log event
- Acceptable types of log messages: ERROR, WARNING, CRITICAL, DEBUG, INFO

Value	Severity	Keyword	Description	Condition
0	Emergency	emerg	System is unusable	A panic condition
1	Alert	alert	Action must be taken immediately	A condition that should be corrected immediately, such as a corrupted system database
2	Critical	crit	Critical conditions	Hard device errors
3	Error	err	Error conditions	
4	Warning	warn	Warning conditions	
5	Notice	notice	Normal but significant conditions	Conditions that are not errors, but that may require special handling
6	Informational	info	Informational messages	
7	Debug	debug	Debug level messages	Messages that contain information normally of use only when debugging a program

- Errors must be easily identifiable and machine readable in the log file
- Format example:  
YYYYMMDD hh:mm:ss, LEVEL, CODE\_MODULE, MESSAGE

## Monitoring

NCO will need the following from the DevOrg in order to set up the monitoring:

- Deliver what should be monitored, either the output or the processes, not both
  - Include when the data or processing is late
  - Include what steps to take if
- List of critical processes that comprise the application that should be monitored
  - Include if the processes runs all the time, or if it has a cadence of starting/stopping
  - A few sentences of documentation notes of what each process does
- List of web sites or other external points of access (sites required to be available to the public)
- Log location on disk
  - Error message pattern to search for in the log
- Nominal disk usage for local files, excluding shared data files
- The verbosity level of the log files must be adjustable through configuration files
- Two points of contact from the DevOrg who will provide Tier 3 support, on an as needed basis

NCO will monitor the execution of the applications on the IDP system using the Big Brother system. In the event of a problem, the 24x7 staff will have contact information for the on-call personnel in the Big Brother information pages. The on-call personnel will troubleshoot the issue and will, if necessary, contact the DevOrg for assistance.

## Application Input Failures

If an application relies on data-of-opportunity as input, it must react to the loss of such data without catastrophic failures or any intervention to recover.

## Protecting IDP Infrastructure and Abusive Users

All software running on IDP must have all public facing components clearly identified so that parts of the application vulnerable to abusive users can be isolated and monitored. Taking action to block abusive users requires the coordination and collaboration of the OBT and IWSB and possibly NSB. The DevOrg shall design the application so that it is robust against abusive users and the internal actions of residing on a shared infrastructure.

## Software Syntax Standards

This section applies to variables applications must use. Such as the home location for ldm, psql, etc.,:

- Keep variable names consistent when using capitals and underscores

- variables must be descriptive and contain one or more words
- One character variable names must only be used in loops or for temporary variables
- Prefix all names with “NCEP\_” for example NCEP\_LDM\_HOME would have the home location for the LDM codelog
- Line length must not exceed 80 or 100 characters with no more than one statement per line
- Arrays must be constructed in a dynamic data structure, never static

## Documentation Standards

### Templates

Standardized templates have been established and adopted by the OBT for all project artifacts. Key among those are templates for project Requirements, Requirements Traceability Verification Matrix (RTVM), Test Cases, and project Build Specifications. These templates shall be maintained in a standard location on Google Drive in folders for each functional area. During the kickoff meeting the applications specifications will be acquired and stored in the Build Spec Sheet Template. For clarification, please contact the OBT Lead.

### Code Management

Application releases are to be version tagged based on a clearly defined methodology that is agreed between the Development teams and the IDP Onboarding team and must follow these standards:

[https://www.nco.ncep.noaa.gov/idsb/Version\\_Numbering\\_Standard.pdf](https://www.nco.ncep.noaa.gov/idsb/Version_Numbering_Standard.pdf)

Git is the mandated Supply Chain Management (SCM) platform. A VLAB/Redmine project must be created for every project that is onboarded, and appropriate Git repositories are to be tied with these projects

- A document that outlines how the application is laid out in the SCM system shall be provided by the Development Organization. Any changes made to the software code structure shall be accurately captured in updates to this document.

### Code Release Documents

When the developer is ready to hand over code the following documents will be included:

- Installation instructions

- Testing verification steps
- Internal-only technical release notes
- Public release notes
- Resolution of all Redmine tickets
- Written assertion that full end to end testing was completed prior to code handoff

## Installation Instructions

The developer must provide an Installation Document for each version release. This document must include at the very least the following:

- Step by step instructions on how to build the application. The build process must be automated using a script and instructions provided on how to run the script.
- Step by step instructions on how to install the application from scratch or as an upgrade to application depending on the type of release it is (new version vs. small upgrade).
  - Where it is possible, an install script will be provided
  - If the process calls for an install on several numerical systems, all systems must be listed explicitly. Do no shorthand the instructions to “vm-lnx-example1-40”, a list of all VMs must be written out.
- Step by step instructions on how to back out the changes.
- Instructions on how to verify a successful installation must be provided for every step
- Identify any environmental files that need to be sourced

## Testing and Validation Standards

### End to End testing

End to end testing with all upstream and downstream partners is a vital check to ensure that prior to installing in production, all partners have verified their ability to provide, receive, or process new data.

Development organizations need to keep track of and provide at delivery time any modifications to the following:

- specific partners, including contact information
- the data delivery or retrieval method

Onboarding will provide a testing platform to use which mimics production workflow.

## Load Testing

Automated programs shall be provided that are designed to test the application under the maximum theoretical loads. When determining the maximum theoretical load, weather conditions at all times of year, and under all circumstances, should be considered as factors. For non-weather based systems, such as emergency and telecommunications systems, a worst case scenario should be envisioned in determining maximum theoretical system load.

When code is delivered to the onboarding team the testing plan must include:

- Endpoints to test the load against
- Load tests for both average days and maximum weather days
- Response time thresholds for all tests
- Average and maximum number of hits expected

## Installation Scripts

The software Installation process must be config file driven. A script setting up the required environmental variables and what their value should be for the different configurations should be centralized to only be changed in one location. This will allow seamless installations between dev, qa, and ops.

## Verification Steps

Verification can and should be viewed from different angles. It is important to confirm that the customer experience is what they expect, regardless of how the application is implemented, but knowledge of implementation details allows for additional tests that would not be obvious simply from a description of the application's capabilities (for example, if there is an internal message queuing system, you might see if it can fill up under extreme circumstances). Functioning of the system is obviously important, but it is also useful to know how the system breaks for various types of failures, to ensure that monitoring and debug procedures are adequate. In addition, various operations such as software installation itself must be verified.

At a minimum, the developers must provide tests to confirm that the application meets all functional requirements from a customer perspective. Each test should have the "action/input" to be done, and the "results expected", clearly outlined.

Development Organizations shall implement appropriate Unit Test Harnesses to implement maximum code coverage in their test cases. A supporting artifact/document that captures the results of these tests are required for review by the Onboarding Team.

Working with the IDP OBT, the development organization will provide tests designed with a specific knowledge of the implementation details of the application, such as:

1. Optimal queue sizes
2. Maximum data rates
3. Maximum down time for failovers and upgrades

The OBT will perform appropriate trample tests. These do not check proper function during normal conditions, but instead characterize how the application fails under untenable conditions (disks filling up, network unavailable, etc.). Our experience has been that some types of failures, for example a mount disappearing, can prove extremely difficult to spot and debug even by primary support personnel during normal work hours, and discoveries made during these tests have been the motivation for new monitoring tests and new debug procedures.

## Verifying Failovers

It is required that specific verification steps and instructions be provided so that an application can be ensured to be working properly after a failover. Verification steps must be automated by using a script or other executable program.

## DataFlow

Applications to be onboarded onto IDP will likely have input data requirements as well as output generated by the application. This section focuses on what data is required, how new data can be retrieved, and how application generated output can be made available.

### Input

All operational applications shall receive inputs from operational systems and never from QA or development systems.

All input datasets used by the application must be described, including:

- Type of data (*model, observation, etc*)
- Source (*NOAA, other government agencies, commercial, etc*)
- Source Type (*FTP, HTTP, LDM, etc*)
- Update Frequency (*hourly, daily, etc*)
- Impact of data being unavailable (*output degradation, application unable to run, etc*)

- Any restrictions to data (*sensitive data from commercial or foreign government source*)
- Any agreements relevant to the retrieval and use of the data.

Requests for data must be submitted to the Dataflow Team using this request form:

[https://docs.google.com/forms/d/e/1FAIpQLSevecbKGt-T3uTSdRdvsvw-KqtDKadBKH1UG\\_2UHZIM\\_nRjTA/viewform](https://docs.google.com/forms/d/e/1FAIpQLSevecbKGt-T3uTSdRdvsvw-KqtDKadBKH1UG_2UHZIM_nRjTA/viewform)

## Output

All output datasets created by the application must be described, including:

- Type and Format of data (*Radar GRIB2, Satellite NetCDF, etc*)
- Individual and aggregate size of data (*ex: Between 1 - 10MB each, keeping 1000 files, disk usage up to 10 GB*)
- Dissemination platforms (*FTP Server, LDM, http service, etc*)
- Requirement for amount of data retention. If the application requires an archive, there needs to be a resource commitment to buy disk for the archive.
- Restrictions on distribution and use of the output data.

## User Data Access

All external users who need to authenticate/authorize into the application should use the enterprise solution - RedHat SSO. This software package is in beta right now, please reach out to OBT if this appears to be a solution that satisfies any application requirements.

## Configuration Management

NCO follows very strict guidelines with all operational changes. Please review the configuration management documents for reference:

<https://sites.google.com/a/noaa.gov/nws-ncep-nco-cm/?pli=1>



## Appendix I - Summary

### Hardware

- Have a hardware specification for each Virtual Machine (VM)
- Red Hat Enterprise Linux: Version 7.7
- Requirement for any local disk space
- Requirement for any shared NFS mounts and local mounts

### Software

- App system and configuration must be the same between sites
- Do not embed code
- Scripts:
  - Should be identical (with some exceptions see above)
  - Identify shell being used in first line
  - Echo errors to logs
  - Make readable via blank lines and many comments
  - Permissions set to 775
- Supported languages: C++, Java, Perl, JavaScript, HTML5, Bash, Python (Anaconda)
- Web Servers supported: Apache Version 2.4 and TomCat Version 7
- Queuing mechanisms allowed: Local Data Manager (LDM) Version 6.13 and RabbitMQ
- Use only the allowed databases, compilers/ IDEs, and build tools (see above for specifications)
- All 3rd party software must be identified by development organization and approved by the Director of NCO

### Failover

- 15 minutes or less (<1 min for WWA apps)
- No data loss and limited data latency
- Verification within 10 minutes after failover
- Limited human input required

### Logging

- Log format example YYYYMMDD hh:mm:ss, LEVEL, CODE\_MODULE, MESSAGE
- "log\_cleanup" script will reside under \$HOME/config called
- 7 day retention
- Zip logs older than a day
- The level of logging must be able to be changed in real time
- Logs will be written to the dedicated NetApp or local application volume; never under the local root volume

### Monitoring

- Dev will provide:
  - List of processes, web sites and log file names and locations
  - Keywords to look for error messages in the logs
  - Expected disk usage
  - Two points of contact from the DevOrg who will provide Tier 3 support

### Application Input Failures

- If an application depends on certain data as input...
  - Loss of data will not crash the application
  - Application will recover when data starts flowing again

### Protecting IDP Infrastructure and Abusive Users

- Portions of the application which are public facing will be identified so they can be protected from abuse

### Software Syntax Standards

- Coding standards - see section

### Documentation Standards

- Templates - must use them
- Code management :  
[https://www.nco.ncep.noaa.gov/idsb/Version\\_Numbering\\_Standard.pdf](https://www.nco.ncep.noaa.gov/idsb/Version_Numbering_Standard.pdf)
- Code release documents
- Installation instructions

### Testing and Validation Standards

- Load testing
- 

### DataFlow

- All input datasets used by the application must be described
  - Type of data / source / source type / update frequency
  - Impact if data is unavailable
  - Access restrictions
- All output datasets generated by the application must be described
  - Type and format
  - Size of data
  - Dissemination platform
  - Data retention requirements
  - Restrictions on distribution and use of the output data

## Configuration Management

- Guidelines: <https://sites.google.com/a/noaa.gov/nws-ncep-nco-cm/?pli=1>