Lustre* Triage & Support

Legal Disclaimer

- THIS DOCUMENT AND RELATED MATERIALS AND INFORMATION ARE PROVIDED "AS IS" WITH NO WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION, OR SAMPLE. INTEL ASSUMES NO RESPONSIBILITY FOR ANY ERRORS CONTAINED IN THIS DOCUMENT AND HAS NO LIABILITIES OR OBLIGATIONS FOR ANY DAMAGES ARISING FROM OR IN CONNECTION WITH THE USE OF THIS DOCUMENT.
- All products, product descriptions, plans, dates, and figures are preliminary based on current expectations and subject to change without notice. Availability in different channels may vary.
- Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation in the United States and other countries.
- *Other names and brands may be claimed as the property of others.
-

Overview

Module Overview

You are the support person - you will be called on for support

- What do you need to do?
- What type of support agreement do you have with Intel's High Performance Data Division (HPDD)?

Some of your tasks may include:

- Collecting information
- Determining whether you need to resolve the issue, or should Intel be engaged?
- Relay information about the issue to Intel
- Track the issue as it gets resolved
- Serve as an intermediary between Intel and your customer throughout the lifecycle of the issue

This module teaches you how to accomplish these steps

Tracking Lustre* Issues

When a qualified Reseller lands its first support PO, a private project is first created for that partner within the Intel Online Issue - Tracking Database known as the HPDD JIRA system (**https://jira.hpdd.intel.com/secure/Dashboard.jspa**).

A reseller can also signup for a private JIRA project by visiting the Support landing page at the Intel Lustre* Reseller Portal (**http://lustre.intel.com**).

If your organization already has a private JIRA project you will be able to see it, provided you have been registered as a member of the project.

If you are not registered, or if your organization does not yet have a dedicated private project, then you will only be able to see the community support project know as Lustre* (LU).

Triage - Who Owns the Issue

Intel Support Levels & SLAs

Support for Intel's Lustre* products falls into the following categories:

- Level 0 – this is the level of support that is provided by all Resellers (partners) to their end customers, as the foundation of the support relationship.
  - The most important element of Level 0 support is to maintain contact with the end customer throughout the support engagement.
- Level 1, Level 2 – these levels of support are often provided by Intel, but may be provided by suitably qualified partners who have successfully completed product and support training.
- Level 3 – this level of support is always provided by Intel.
  - Level 3 support may involve making changes to Lustre* code.
  - Level 3 support requires in depth knowledge of Lustre* code internals.

Responsibilities - Level 0 Support

Level 0 | Level 1 | Level 2 | Level 3

All Reseller Partners are expected to provide Level 0 support to their end customers.

Partner responsibilities for Level 0:

- Receive and handle initial contact from end customer
- Validate that the issue is related to Intel® EE for Lustre* Software or one of its components (Intel® Manager for Lustre* software, HSM, HAL, HAM)
- Check that a current support contract with Intel exists for this customer
- Create the basic JIRA ticket on behalf of the end customer
- Manage the relationship with the end customer during resolution of the issue, until the issue is resolved.

Responsibilities - Level 1 Support

Level 0 | **Level 1** | Level 2 | Level 3

Problem and/or bug validation

Provide relevant information about the End Users environment, setup and configuration

Review solutions database for known problems with a known resolution

Research to determine whether the problem already exists in the HPDD JIRA Issue Tracking Database

- Note: A partner has additional visibility to their own private project only, and not tickets from any other partners private projects

Logging of validated problems in a JIRA ticket

Ensure the validated problem report is complete

- Fully describe the problem, its impact
- Provide available diagnostic and logging information

Follow-up interactions along with best efforts to resolve the problem

Responsibilities - Level 2 Support

Level 0 | Level 1 | **Level 2** | Level 3

Make commercially reasonable efforts to:

- Reproduce and diagnose the problem
- Resolve or reduce severity of problems by introducing:
  - Known work-arounds
  - Patch(es)
  - Code corrections
  - Operating System and Software upgrades, as recommended

Validate whether any fixes provided resolve the problem

Communicate steps, action taken and results in the JIRA ticket

- Include any code changes and testing results
- Follow-up interactions, along with best efforts to resolve the problem

Responsibilities - Level 3 Support

Level 0 | Level 1 | Level 2 | **Level 3**

Just to reiterate – Intel is always the provider of Level 3 support, regardless of how Level 1 and Level 2 support is provided.

Supply successful problem resolution to reproducible problems

Validate fixes made via regression testing

Develop new test cases, where applicable

- As a result of a root cause determination

Communicate steps of action or resolution and code changes

- Includes testing results in Intel's online issue tracking database

One or more of the following:

- Submit final resolution to the Lustre* master tree
- Issue a hotfix
- Perform updates to the Intel software code base

Support SLA – Response Times

Severity Level	Severity 1	Severity 2	Severity 3	Severity 4
Initial Response	Maximum response time 4 hours	3 Days	5 Days	As available
Work Schedule	Continuous	Workday	Workday	As available

Severity 1 - Major Impact

- Continued file system outage affecting a large set of End Users

Severity 2 – High Impact

- File system or cluster is down and non-functional
- Problem causing sporadic or partial cluster outages
- Adverse impact to End User interaction with the file system

Severity 3 - Medium Impact

- Problem that includes partial non-critical loss of cluster access
- Impairs some operations but allows the End User to continue with workarounds

Severity 4 – Minor Issues

- Minor issue with limited or no loss in functionality
- May also be used for recommendations for future product enhancements or modifications

Issue Reporting

Tracking System — Overview

What type of information may be needed?

- Name of site experiencing the issue
- Name of person/company submitting the new issue
- Type of request - bug, new feature, etc.
- Description of the request, including:
  - Bug: Evidence of, and how to reproduce the problem
  - New Feature: Justification
- Severity of the impact (if applicable)
- System and software information - affected version(s)

Tracking Systems — Access

HPDD JIRA Issue Tracking Database

- Information accessible w/o an account

Reseller Partner's 'private' Project - Requires Intel support agreement

- Must be authenticated to access information
- Who can view issues you create?
  - HPDD Engineering
  - You, and others in your company

Tracking Systems — Accounts

HPDD JIRA Issue Tracking Database

- You will need to create your own JIRA account at jira.hpdd.intel.com

Partner Project within the HPDD JIRA Issue Tracking Database

- Intel will create and provision your account
  - With appropriate access and authorization
  - Necessary customizations for your company
  - Then you can register here:
    https://jira.hpdd.intel.com/secure/Dashboard.jspa

Gathering Information

Collecting Information - How?

Your first introduction to the problem will be:

- Via your company's interface to your customer

How will you communicate?

- Phone, email?

Either way, **collecting information is a process**

The process should be:

- Get your checklist
- Get the customer to explain the problem
- Get the information **needed** to create a JIRA ticket
- Get any additional information that should be collected

Collecting Information - Sample Checklist

Recall from earlier, the information that you need to know:

- Project - Your company
- Issue Type - Bug, New Feature
- Summary - Title for the JIRA ticket
- Priority - How critical you (not your customer) think it is
- Environment - Customer site info, including software versions
- Description - Describe the problem as best you can
- Attachments - Files that might be provided by your customer; e.g., screenshots, logs, chroma-diagnostics

Collecting Information - Process

Defining the process may take a few iterations

The process will likely include:

- Establish a method for collecting information on the checklist; e.g.
  - Web interface, printed document, etc.
- Phone or email to collect the information
- Analyze the information - look for missing aspects
- Open a JIRA support ticket using the information

Collecting Information – What

Lustre* related issues

You will be looking for evidence of a problem with the customer's Lustre* file system

Your customer should look for recent messages on the console (dmesg) or in the system logs (/var/log/messages) on the client. This information can be attached to a JIRA ticket.

Examples:

```
OSS dmesg output:
LustreError: 49269:0:(events.c:447:server_bulk_callback()) event type 3, status -103, desc
ffff883a20842000

LNetError: 49269:0:(o2iblnd_cb.c:3036:kiblnd_check_txs_locked()) Timed out tx: active_txs,
12 seconds
```

Lustre* Manual: **https://wiki.hpdd.intel.com/display/PUB/Documentation**

Note that the absence of Lustre* errors is quite good evidence that the problem lies elsewhere.

Collecting Information – What (cont.)

Problems with Intel® Manager for Lustre* software

Intel® Manager for Lustre* software is a user-friendly front end GUI that allows you to set up an EE Lustre* file system and then manage it.

Intel® Manager for Lustre* software will report issues with the file system using a powerful

The best way to help your end customers with Intel® Manager for Lustre* software is to make sure they have a copy of the Partner Installation Guide and the Intel® Manager for Lustre* software User Manual.

Documentation can be found at the 'Support' page in the Reseller Portal here – **http://lustre.intel.com**

Collecting Information – What (cont.)

Problems with Intel® Manager for Lustre* software – obtaining chroma-diagnostics

- Log onto each server of the problem system (MGS, OSS, MDS)
- At the command prompt type: chroma-diagnostics
- Locate the logs in /var/log

Upload logs and traces to JIRA ticket:

```
[root@mds1 ~]# chroma-diagnostics
Collecting diagnostic files
Detected devices
Listed installed packages
Listed cibadmin --query
Listed: pcs config show
Listed: crm_mon -1r
Finger printed Intel Manager for Lustre® installation
Listed running processes
listed PCI devices
listed file system disk space.
....
....
Copied 1 log files.
Diagnostic collection is completed.
/var/log/diagnostics_20140909T210806_st76-mds1.tar.lzma
```

Example: Creating a New Issue in the Lustre* HPDD JIRA

Tracking System — Authenticated

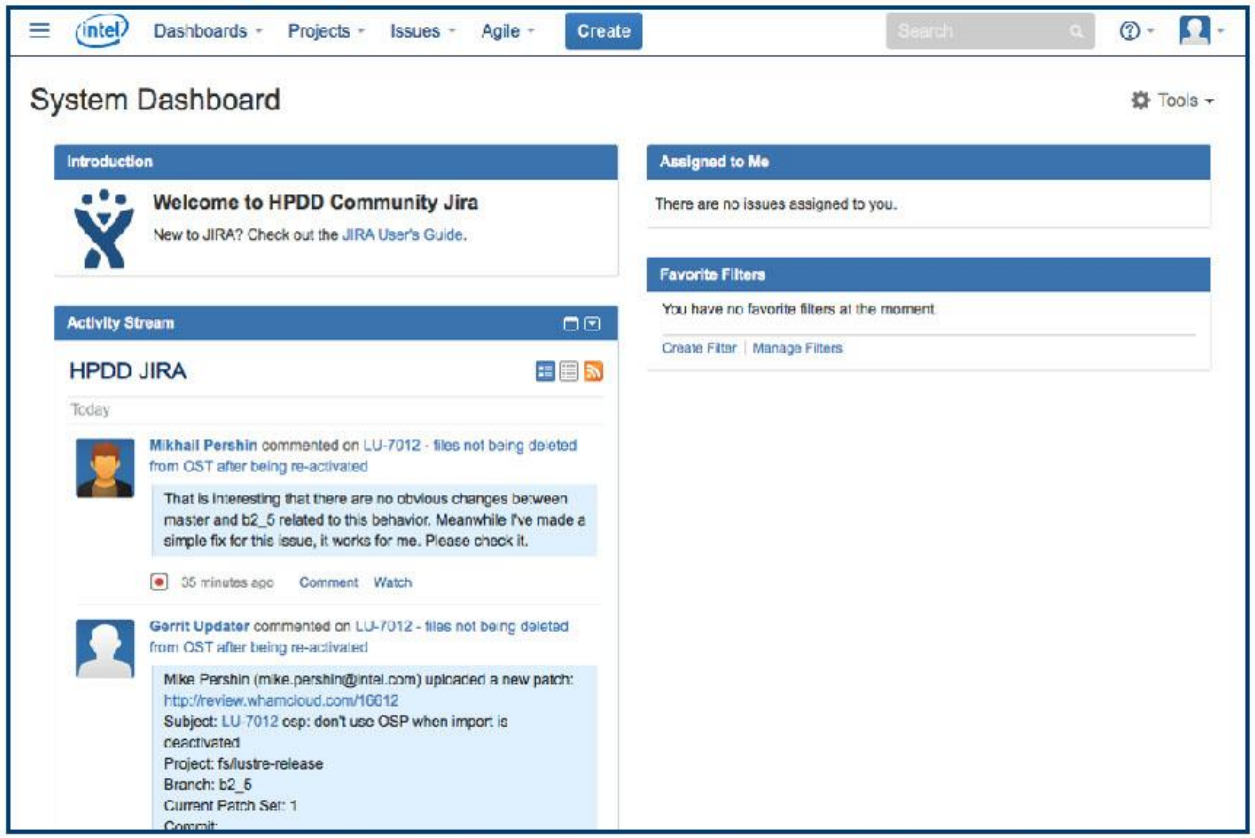

Tracking System — Creating a JIRA

Begin by clicking the “Create” button at the top - A new dialog opens

- **Important note: JIRA is ‘sticky’**
  - *This means that when you create a new issue, the default Project will be the project you **last used**. Make sure you have selected your private JIRA Project.*

The dialog has several fields to enter information

- You will not need to complete every field

Recommended fields to provide information are:

- Project, Issue Type, Summary, Priority, Environment, Description
- Uploading supporting documents such as screenshots, logs and chroma-diagnostics via the Attachment button
  - Note: 30MB is the maximum file size. For attachments larger than 30MB, refer to the Anonymous Write-Only FTP Upload Service

Tracking System — Create Issue

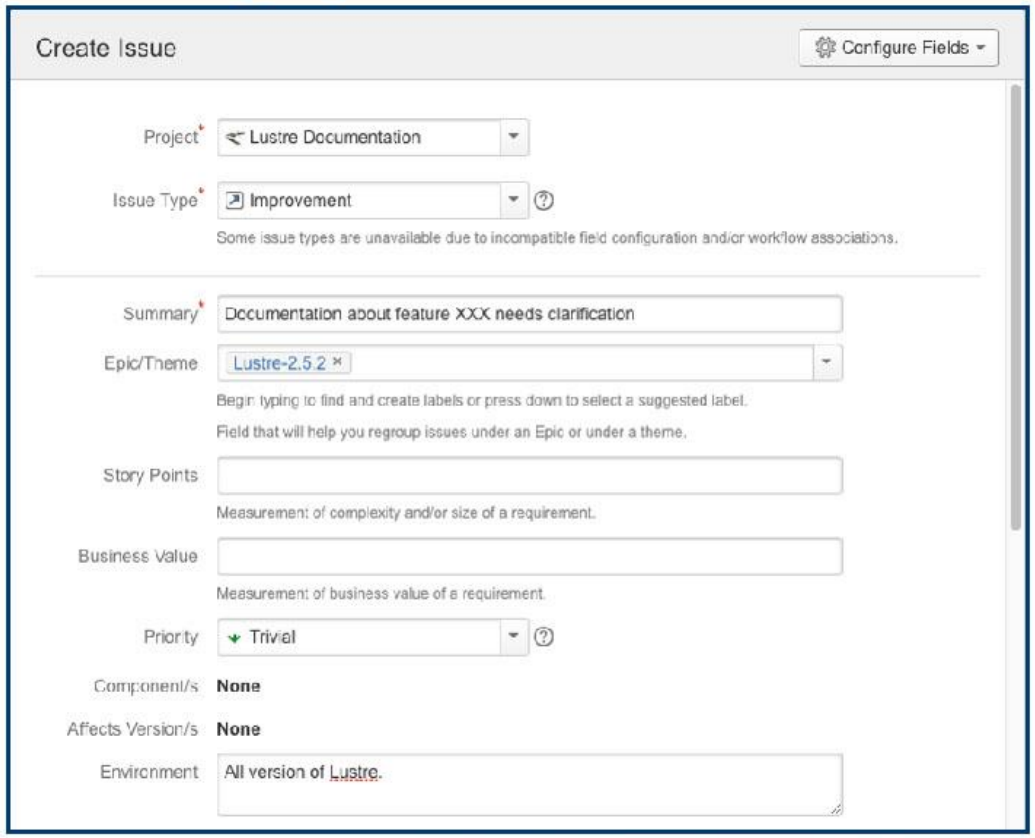

Tracking System — Create Issue (cont.)

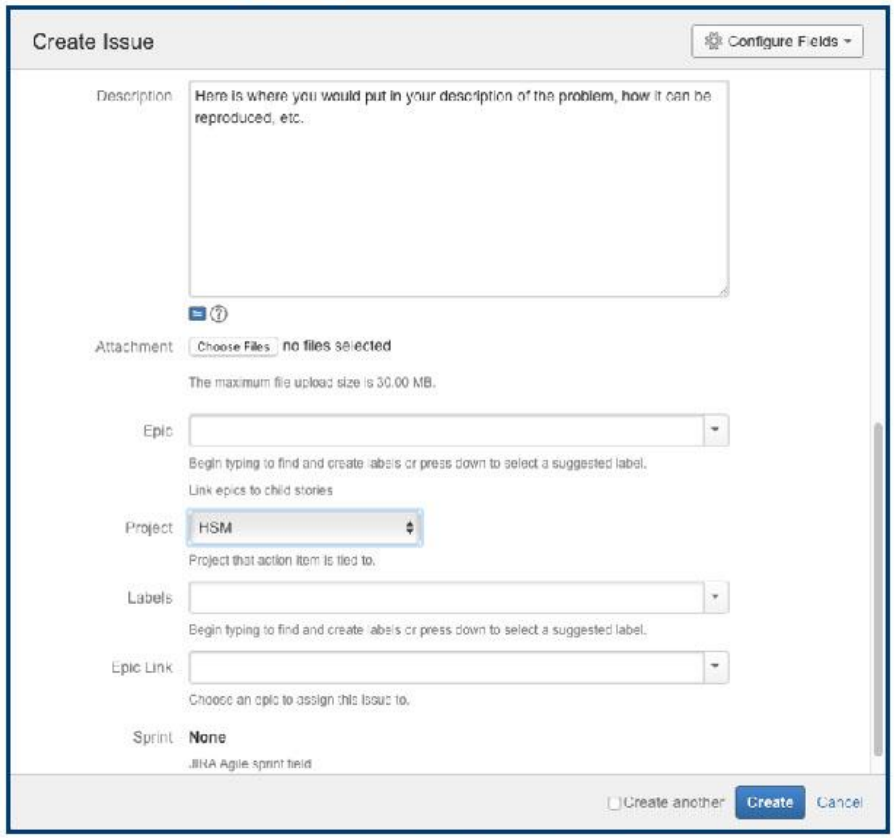

Tracking System — Help Dialog

Issue Types

JIRA can be used to track many different types of issues. The currently defined issue types are listed below. In addition, you can add more in the administration section.

FOR REGULAR ISSUES

Bug	A problem which impairs or prevents the functions of the product.
Epic	A big user story that needs to be broken down.
Improvement	An improvement or enhancement to an existing feature or task.
New Feature	A new feature of the product, which has yet to be developed.
Question/Request	A product question or request
Story	An user story
Task	A task that needs to be done.

FOR SUB-TASK ISSUES

Review task	A document or code review that needs to be done.
Sub-Task	The sub-task of the issue
Technical task	A technical task.

Priority Levels

An issue has a priority level which indicates its importance. The currently defined priorities are listed below. In addition, you can add more priority levels in the administration section.

Blocker	Blocks development and/or testing work, production could not run.
Critical	Crashes, loss of data, severe memory leak.
Major	Major loss of function.
Minor	Minor loss of function, or other problem where easy workaround is present.
Trivial	Cosmetic problem like misspelt words or misaligned text.

Not shown on slide:

- Statuses
  - Status categories
  - Issue statuses
- Resolutions

Anonymous Write-Only FTP Upload Service

ftp.hpdd.intel.com

- Use the 'uploads' directory and create a new subdirectory using your JIRA issue as a name.

In the following example, there are three debug logs in a single directory and the JIRA issue LU-4242 has been created. After completing the upload, please update the relevant issue with a note mentioning the upload, so that our engineers know where to find your logs.

```
$ ncftp ftp.hpdd.intel.com
Logging in...
Login successful.
Logged in to ftp.hpdd.intel.com.
ncftp / > cd uploads
Directory successfully changed.
ncftp /uploads > mkdir LU-4242
ncftp /uploads > cd LU-4242
Directory successfully changed.
ncftp /uploads/LU-4242 > put *
mds-debug: 97.66 MB 11.22 MB/s
oss-00-debug: 117.19 MB 11.16 MB/s
oss-01-debug: 117.48 MB 11.18 MB/s
```

Tracking Information

Tracking Your JIRA

After requests are submitted, they are triaged:

- Assigned a priority within Intel
- Assigned to an Engineer

These actions will trigger an email to the reporting Engineer - You

Each additional update to the ticket will also trigger emails to you

- In this way, you will be able to be informed of the current status
- Significant changes may be fed back to your customer

In each email, there is a link to the ticket

- The link provides a way to review the ticket whenever you like

Summary

Summary

- The HPDD JIRA Issue Tracking Database uses Atlassian's JIRA product to track issues
- Partner Projects within the HPDD JIRA Issue Tracking Database will be provided credentials
- To create new issues on behalf of your company, you will:
  - Collect information from your customer
  - Log in to the HPDD JIRA Issue Tracking Database
  - Click the Create button at the top of the page
  - Complete the appropriate fields
  - Click the Create button at the bottom of the dialog
- Check your email, and verify the details of the new issue

Features and Benefits of Intel HPDD Partner Support

Features	Benefits
The Intel HPDD team includes the most experienced Lustre* engineers available, many of whom have been involved in the development and support of Lustre* since its inception.	HPDD Intel EE Reseller Partners receive the highest level of expert technical support available anywhere.
HPDD is a worldwide distributed organization – we have support engineers in ten time zones worldwide.	HPDD support 'never sleeps' and is always available to respond to customer issues as they are reported, within the contractual SLA response times.
Lustre* and IML releases are tightly controlled for quality. Every patch is subject to rigorous expert review. Intensive and continuous 'autotest' processes enable regressions to be promptly identified and fixed.	Even though Lustre* is extremely large and complex, with several hundred thousands of lines of code, it is a high quality product and is stable with most workloads.
HPDD supports hundreds of customers either directly or through its network of partners. This gives us extensive experience across a wide variety of different workloads and environments.	The number of supported end customers is small enough to allow HPDD to maintain close ties, through their Intel EE Reseller Partners, and to know their individual requirements and working environments.
When a SEV 1 customer support event occurs, a JIRA ticket email is broadcast to everyone in HPDD on an hourly basis, and this broadcast is repeated until work has commenced.	HPDD Intel EE Reseller Partners can have the assurance of knowing that they will receive first class support, when they need it most urgently.
HPDD Intel EE Reseller Partners have access to an escalation email alias which goes to members of the support team, and also to HPDD management.	Customers can get immediate high level visibility of their critical support issues within HPDD, which gets them the attention they need.

Congratulations! You have completed:

Lustre* Triage & Support