



IML Roadmap and Development Community

Joe Grund

IML Team Lead

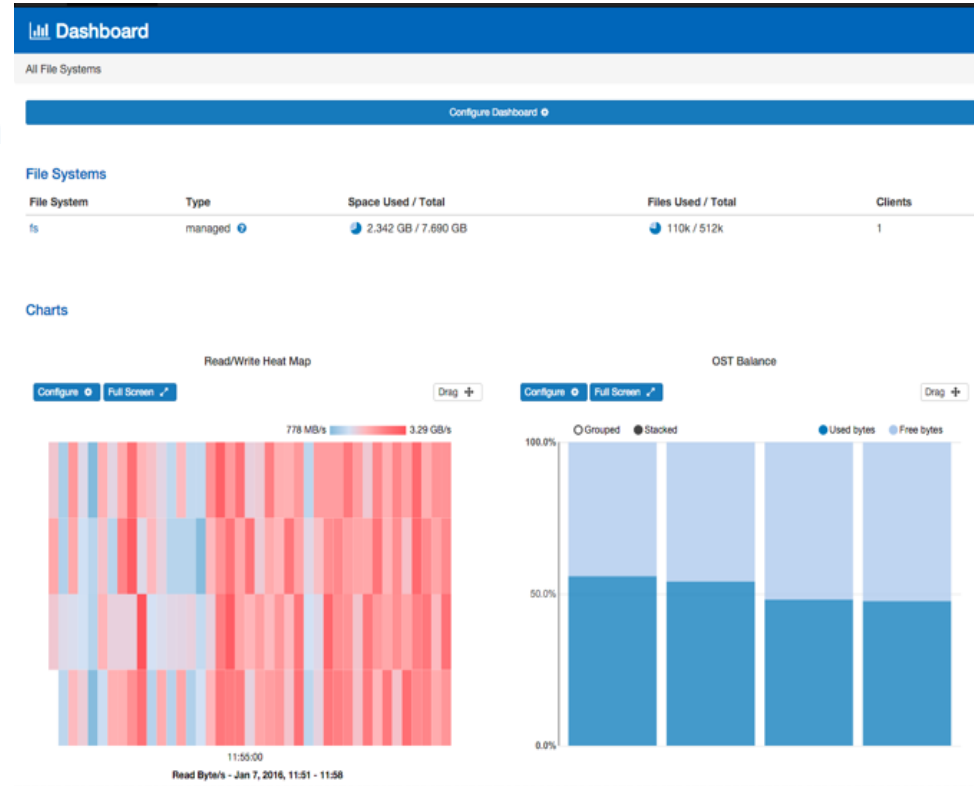
joe.grund@intel.com

Agenda

- IML Background / Overview
- Roadmap
 - 4.1 Release
 - 4.x Series
- Contributing to IML

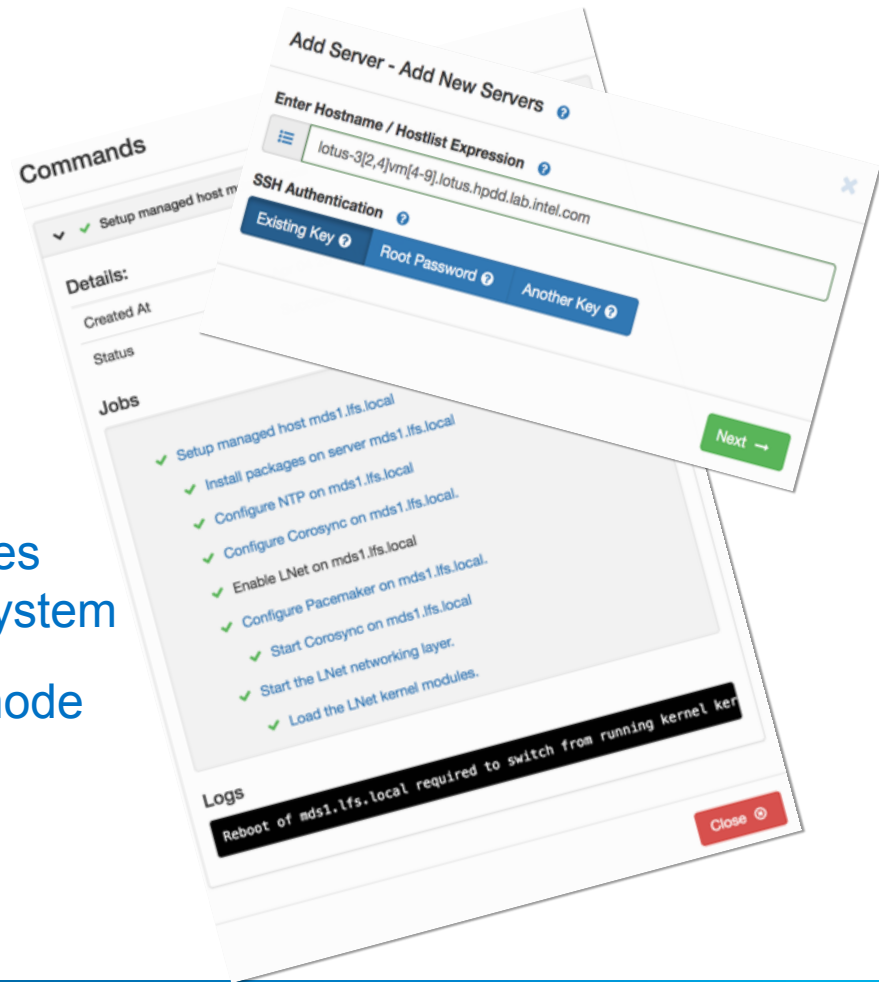
Background

- Intel® Manager for Lustre (IML) is an open source suite of tools for deploying, managing, and monitoring Lustre
- IML simplifies Lustre administration with intuitive interfaces and near real-time feedback
- Works with new and existing Lustre installations
- Monitors performance and system health



Deployment

- Deploy Lustre filesystems from one centralized location
- Near-realtime feedback
- Bring filesystem online from first principles or deploy monitoring for an existing filesystem
- Deploy specialized assets, HSM agent node
- Add more assets over time



Management

- Configure / change state of Lustre and related components
- Uses state-machine to reach end state from different starting points
 - Starting LNet, state machine ensures packages are installed + kernel modules loaded before bringing LNet up
- Handle recovery situations fencing, failover
- Uses Corosync / Pacemaker, with PDUs / IPMI

The screenshot displays the Intel® Management for Lustre™ software interface for a server named 'test003.localdomain'. The interface is organized into several sections:


- Server Detail:** Shows metadata for the server, including Address (test003.localdomain), State (managed), FQDN (test003.localdomain), Nodename (test003), Profile (Managed Storage Server), Boot time (Thursday, January 7, 2016 11:26:08), State changed (Thursday, January 7, 2016 11:26:40), and Alerts (No Issues).
- Pacemaker Configuration:** Shows State (Pacemaker Started) and Alerts (No Issues).
- Corosync Configuration:** Shows Mcast Port (3000), State (Corosync Started), and Alerts (No Issues).
- LNet Detail:** Shows State (LNet Up) and Alerts (No Issues).
- NID Configuration:** A warning message states: "To edit the NID configuration, ensure this host is not a member of an active file system." Below this is a table mapping interfaces to IP addresses, Lustre Network Drivers, and Lustre Networks.

Interface	IP Address	Lustre Network Driver	Lustre Network
eth0	10.0.0.3	tcp	Lustre Network 0
eth2	10.2.0.3	tcp	Lustre Network 2
eth4	10.4.0.3	tcp	Lustre Network 4
ib1	10.1.0.3	ib2ib	Lustre Network 1
ib3	10.3.0.3	ib2ib	Lustre Network 3

Intel® Management for Lustre™ software Build 0 - Copyright © 2015 Intel Corporation. All rights reserved. About Intel® Management for Lustre™ software System Status

Monitoring

- Holistic system metrics
- Rich visualizations
- Drill into filesystem, target, server
- Find and monitor top jobs
- Aggregate logs across cluster
- HSM Copytool activity monitoring
- Alerts to cluster issues
- GUI / Email / API
- Searchable command / event / alert log / history



Job Stats

Select Duration

1 minute 5 minutes 10 minutes

Job	Avg. Read Bandwidth	Min. Read Bandwidth	Max. Read Bandwidth	Avg. Write Bandwidth	Min. Write Bandwidth	Max. Write Bandwidth	Avg. Read IOPS	Min. Read IOPS	Max. Read IOPS	Avg. Write IOPS	Min. Write IOPS	Max. Write IOPS
05.0	218.1 MB/s	82.38 MB/s	333.3 MB/s	185.4 MB/s	58.83 MB/s	328.3 MB/s	20,793	4.3	34	22,891	10.8	34,25
02.0	201.8 MB/s	63.76 MB/s	318.2 MB/s	168.6 MB/s	89.41 MB/s	287.0 MB/s	18,852	5.5	28.65	19,987	9.1	30

Development History

IEEL 1.0 – 2013

- Chroma renamed to IML

IML 2.2 – 2015

- ZFS Monitored mode support
- Enhanced parallel server deployment
- Enhanced command display / drilldown

IML 3.1 - 2016

- ZFS support for managed mode.
- Near realtime jobstats monitoring
- tree-view

Chroma 1.0 – 2012

- Bare Metal install
- Start / Stop / Add / Remove FS and targets
- GUI, REST, CLI interfaces
- HA over Lustre targets

IML 2.1 – 2014

- Heatmap visualization.
- Job stats collection and display
- Hierarchical Storage Management (HSM) support and display
- Custom profile Support
- Near realtime support

IML 3.0 – 2016

- Pacemaker / Corosync config
- NID config
- Queryable status
- Architectural improvements

IML 4.0 - 2017

- First open source release
- Tracks Lustre LTS
- appliance -> services

Roadmap

4.1 Roadmap (Active Development)

- Increase scalability of device detection
- Increase modularity of components
- Start moving towards separate management / monitoring code paths
- Targeting GA Q2 2018

Scope

- libzfs integration
 - <https://github.com/intel-hpdd/intel-manager-for-lustre/issues/535>
- ZED integration
 - <https://github.com/intel-hpdd/intel-manager-for-lustre/issues/536>
- Reactive Architecture
 - <https://github.com/intel-hpdd/intel-manager-for-lustre/issues/533>
- Full Modularity
 - <https://github.com/intel-hpdd/intel-manager-for-lustre/issues/534>

libzfs integration

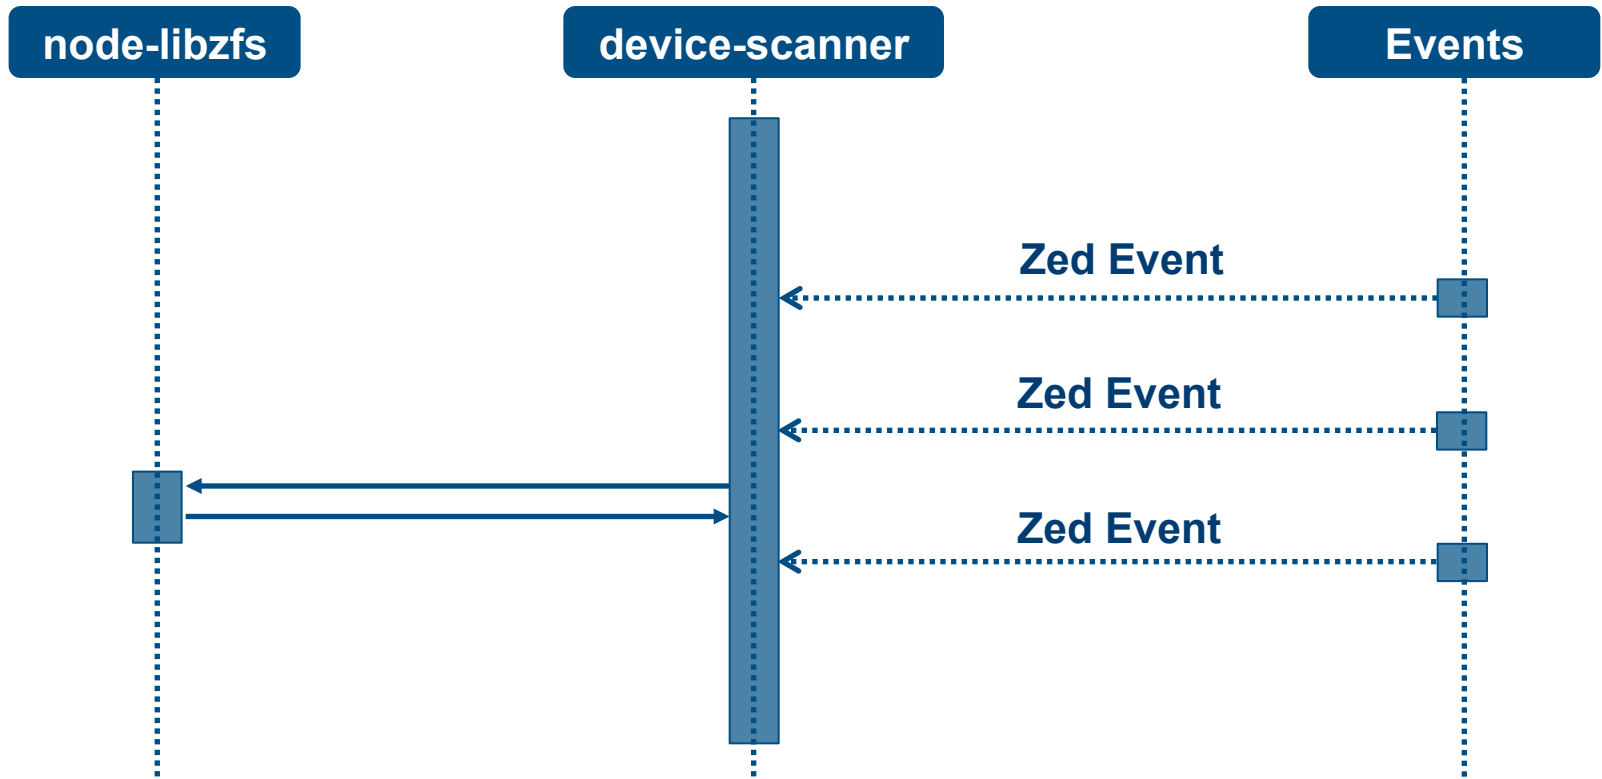
- IML looking to utilize libzfs in 4.1 release for multiple purposes
- Lower level interface over invoking commands / parsing CLI output
- Fine grained collection of pools / datasets / props / VDEV tree
 - easy to collect more data later
- Useful for monitoring + upcoming management enhancements

```
[
{
  "name": "test",
  "guid": "14919184393193585238",
  "health": "ONLINE",
  "hostname": "localhost.localdomain",
  "hostid": 3914625515,
  "state": "ACTIVE",
  "readonly": false,
  "size": 83886080,
  "vdev": {
    "Root": {
      "children": [
        {
          "Mirror": {
            "children": [
              {
                "Disk": {
                  "guid": "0xBE4606AF1C39DC3F",
                  "state": "ONLINE",
                  "path": "/dev/sdb1",
                  "dev_id": "ata-VBOX_HARDDISK_081118FC1221NCJ6G8G1-part1",
                  "phys_path": "pci-0000:00:0d.0-ata-2.0",
                  "whole_disk": true,
                  "is_log": null
                }
              }
            ],
            "is_log": false
          }
        },
        {
          "Disk": {
            "guid": "0xCC43D91716DA2522",
            "state": "ONLINE",
            "path": "/dev/sdc1",
            "dev_id": "ata-VBOX_HARDDISK_081118FC1221NCJ6G8G2-part1",
            "phys_path": "pci-0000:00:0d.0-ata-3.0",
            "whole_disk": true,
            "is_log": null
          }
        }
      ]
    }
  },
  "is_log": false
}
],
]
```

ZED Integration

- IML looking to use ZED in 4.1 release for multiple purposes
 - Discovery of pool / dataset / property / VDEV changes
 - This is currently a manual scanning process
 - Will allow for closer to realtime changes to propagate
 - Better scaling
 - Surfacing alerts in the IML GUI / API
 - Alerting on critical events across a cluster
 - Searchable history of all events across a cluster

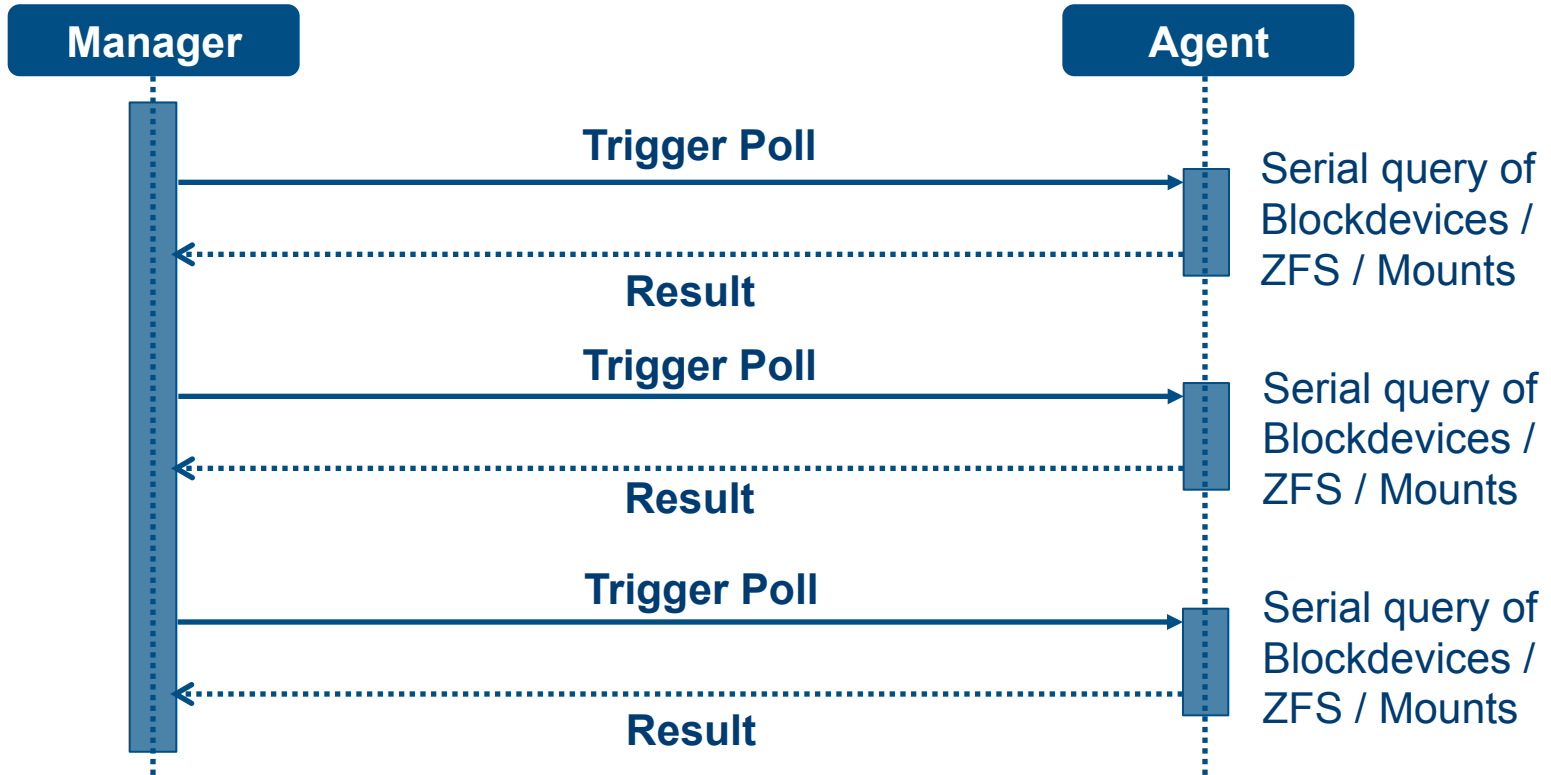
ZED Integration



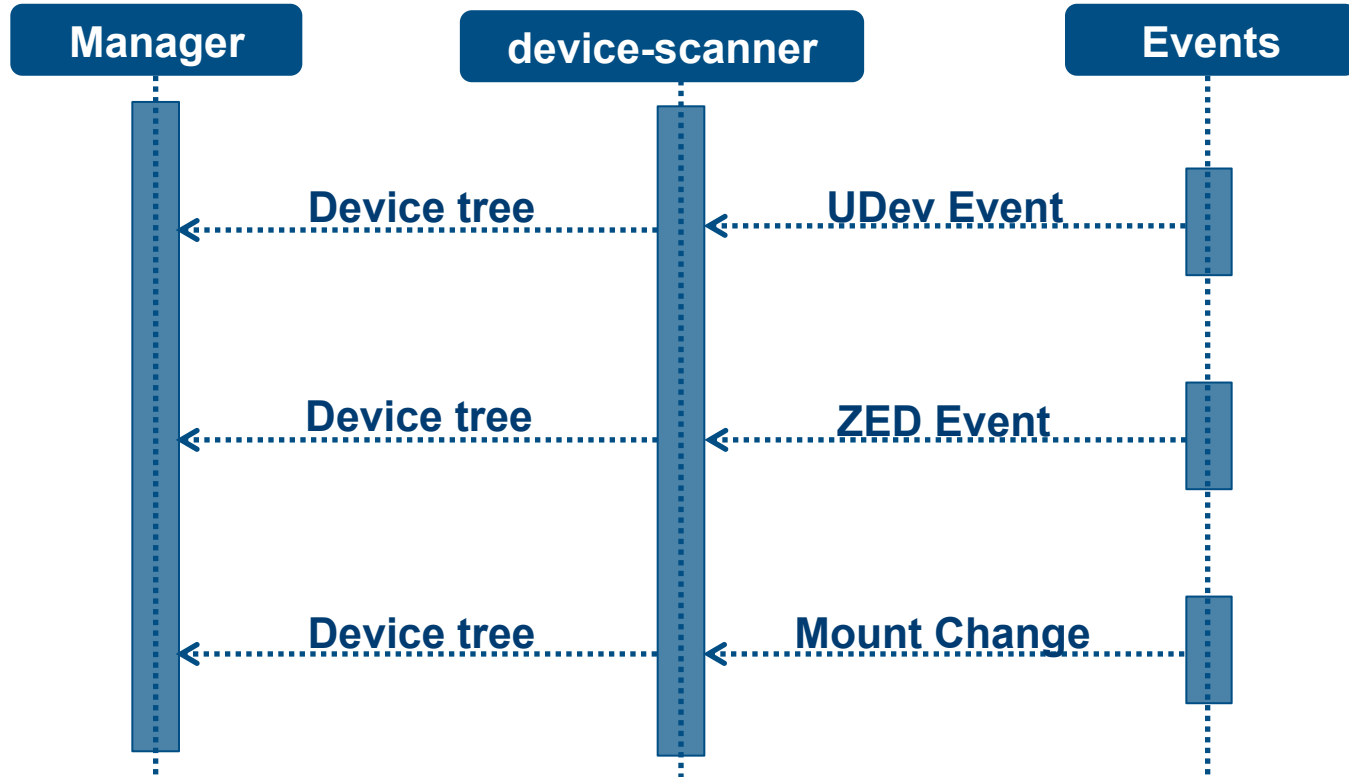
Reactive Architecture

- IML looking to flip device discovery from push to pull
 - Adds further scalability
 - Has lower resource usage
 - More responsive
 - Previous iteration used polling + serial introspection of devices

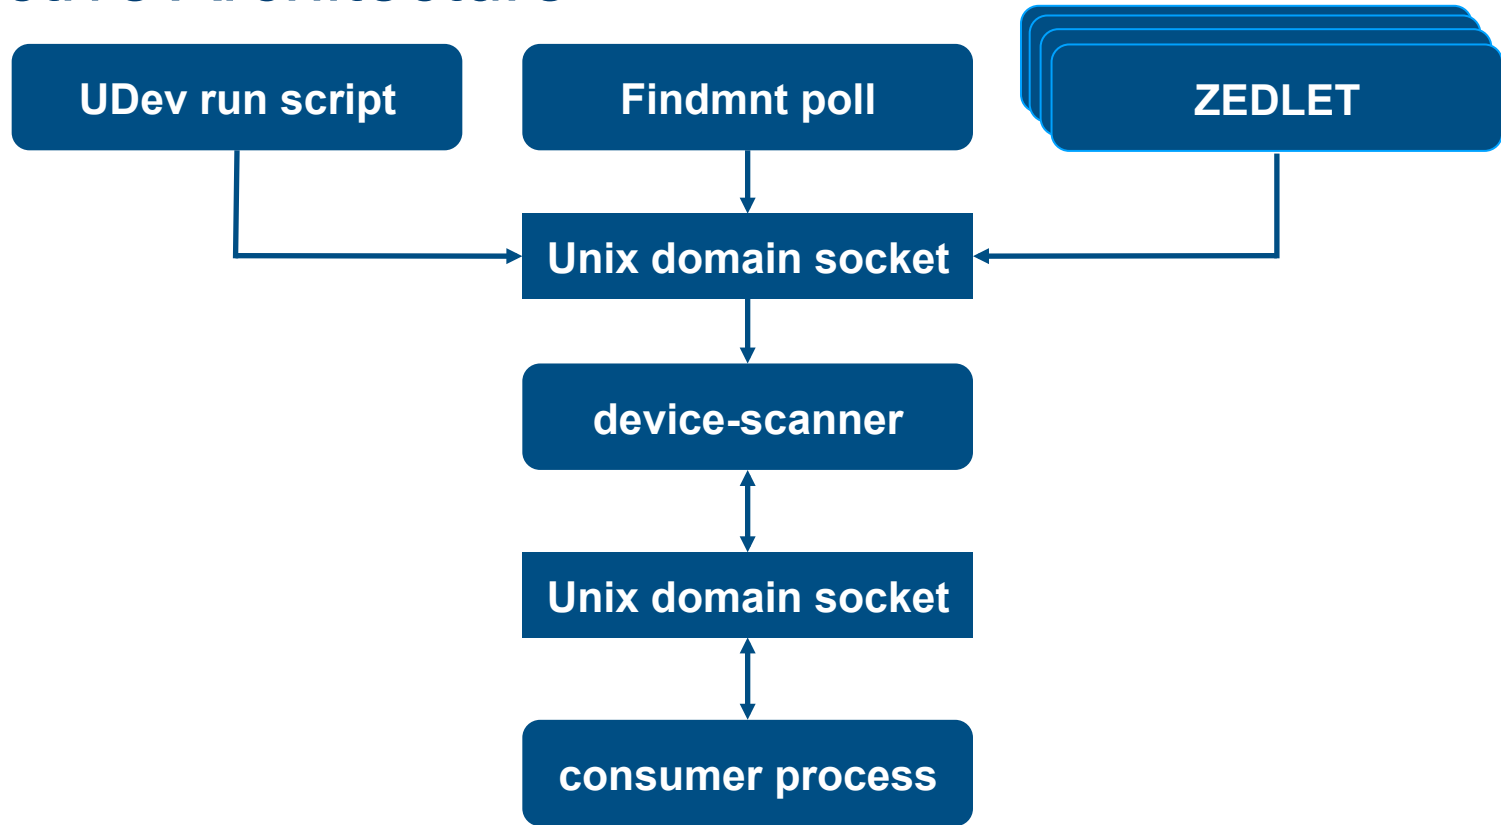
Reactive Architecture



Reactive Architecture



Reactive Architecture



Modularity

- IML looking to deliver itself completely via RPM
 - No more tarball
 - Ship everything via Fedora Copr
- Benefits
 - Semver minor updates via yum update
 - Components evolve independently
- Those not wanting continuous deployment can disable upgrades / perform offline install / upgrade.

4.x Series Roadmap

- Proposed enhancements post 4.1
- Items may be scoped / de-scoped by priority and community feedback
- Prioritize before starting next release
 - IML 4.2 prioritization process to start in coming weeks
- Target quarterly IML enhancement releases
- As needed compatibility + bugfix releases

Enhanced Deployment

- IML should make it even easier to setup Lustre
 - Deploy to large scale clusters with minimal operator intervention
 - Describe ideal cluster state
 - Expose variants as scalable UI widgets
 - Deploy installation in parallel with a single click

Timeseries DB

- IML Should make it easy for users to consume + display stats
 - Use a popular TSDB to collect stats
 - Enable well known API + integration into 3rd party dashboards
 - Better perf / ability to locate stats separately from main db

Use Database Change Events

- IML should use database change events
 - Reduce overhead
 - Push at manager layer
 - Eliminate long-polling, use change events at DB level

Deprecate Storage Plugin

- IML should support more storage arrays as first class citizens
 - Bring functionality in-house for well known plugins
 - Simplify architecture significantly
 - Tightly integrate storage arrays into all of IML

Hardware Data at rest encryption key management

- IML should allow use of hardware encryption across cluster
 - Allow user to encrypt / decrypt all drives or just subset
 - Ability to store pins in a decentralized store
 - Ability to share encrypted pins between nodes for redundancy + HA

Lustre Snapshot Management

IML should be able to manage Lustre snapshots via GUI

- Schedule snapshots for filesystems at some regular interval
- Ad-hoc snapshot on filesystem(s)
- View / delete previous snapshots
- Rollback to a given snapshot
- Rename a snapshot
- Take snapshot at key points (i.e. Lustre upgrade)

Full ZFS Management

- IML should provide full ZFS management
 - Show the list of all pools and datasets
 - Provide drill-down navigation to elicit more detail on a selected target
 - Show the status of pools and datasets
 - Where imported, mounted, error conditions, configuration
 - Management
 - Create zpools / datasets
 - Update volume display to easily setup devices, e.g. RAID-Z
 - JBOD enclosure GUI


I18n Support

- IML should have an easy facility for i18n
 - IML text currently English, but IML is used all over the world
 - Need a way to support non-native English speakers
 - Ideally in a way that allows end-user contribution/translation
 - Modify/contribute *.po files consumed by services
 - Help Wanted!

Contributing to IML

How to get started - Contributing

1. Signup for github.com account if you don't have one
2. Navigate to the repo you are interested in under our organization
<https://github.com/intel-hpdd/>
3. Fork the repo to your personal Github



4. Clone the desired repository locally

```
git clone git@github.com:my-user/GUI.git
```

How to get started - Contributing

5. Set upstream to intel-hpdd repo

```
git remote add upstream git@github.com:intel-hpdd/GUI.git
```

6. Create new working branch

```
git checkout -b my-new-branch
```


7. Make edits, commit + push

```
git commit -s  
git push origin my-new-branch
```

How to get started - Contributing

8. Visit your repo in Github GUI, click pull request button

Your recently pushed branches:

 my-new-branch (less than a minute ago)

 Compare & pull request

9. Checks will run against CI provider

10. Organization members review, request changes or land

Help Wanted

- Check Github issues for help wanted opportunities
 - <https://github.com/issues?utf8=%E2%9C%93&q=is%3Aopen+is%3Aissue+archived%3Afalse+user%3Aintel-hpdd+label%3A%22help+wanted%22>
 - Easy to implement, team guidance
- Open an issue / submit a PR
 - Use a release train model, pull in work once it's done
 - Want your feedback on useful enhancements
 - Milestones are public
 - <https://github.com/intel-hpdd/intel-manager-for-lustre/milestones>

Get in Touch

- Many ways to reach out
 - Join us on gitter.im for any questions
 - <https://gitter.im/intel-hpdd/intel-manager-for-lustre>
 - Request any issues / enhancements on the IML Repo
 - <https://github.com/intel-hpdd/intel-manager-for-lustre/issues/new>
 - Checkout our docs
 - <https://intel-hpdd.github.io/Online-Help/>
 - We have extensive docs on contributing
 - https://intel-hpdd.github.io/Online-Help/docs/Contributor_Docs/cd_TOC.html

Closing

- IML is a project with a long history and continues advancing
 - IML 4.0 was first open source release in Oct 2017
 - Five maintenance updates to 4.0 release
 - IML 4.1 targeted for Q2 2018
 - IML 4.2, next release; prioritization shortly
- We are actively seeking contributors
 - Open tickets or PRs for enhancements / bugfixes
 - Get involved in IML's future!

Legal Notices and Disclaimers

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY RELATING TO SALE AND/OR USE OF INTEL PRODUCTS, INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT, OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life-saving, life-sustaining, critical control or safety systems, or in nuclear facility applications.

Intel products may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

This document may contain information on products in the design phase of development. The information herein is subject to change without notice. Do not finalize a design with this information. Intel may make changes to dates, specifications, product descriptions, and plans referenced in this document at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

No computer system can be absolutely secure.

Intel Corporation or its subsidiaries in the United States and other countries may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

Performance estimates or simulated results based on internal Intel analysis or architecture simulation or modeling are provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice Revision #20110804.

Intel® Advanced Vector Extensions (Intel® AVX)* provides higher throughput to certain processor operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you can learn more at <http://www.intel.com/go/turbo>.

Intel processors of the same SKU may vary in frequency or power as a result of natural variability in the production process.

Intel, the Intel logo, 3D-Xpoint, Optane, Xeon Phi, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

* Other names and brands may be claimed as the property of others.

Copyright © 2018 Intel Corporation. All rights reserved.

