Test infrastructure: Where are we now?

Chris Gearing
April 17th 2013

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

Agenda

1. Intel’s Values
2. What were our aims?
3. What have we done?
4. Where are we going?

Intel’s Values?

Intel’s Values

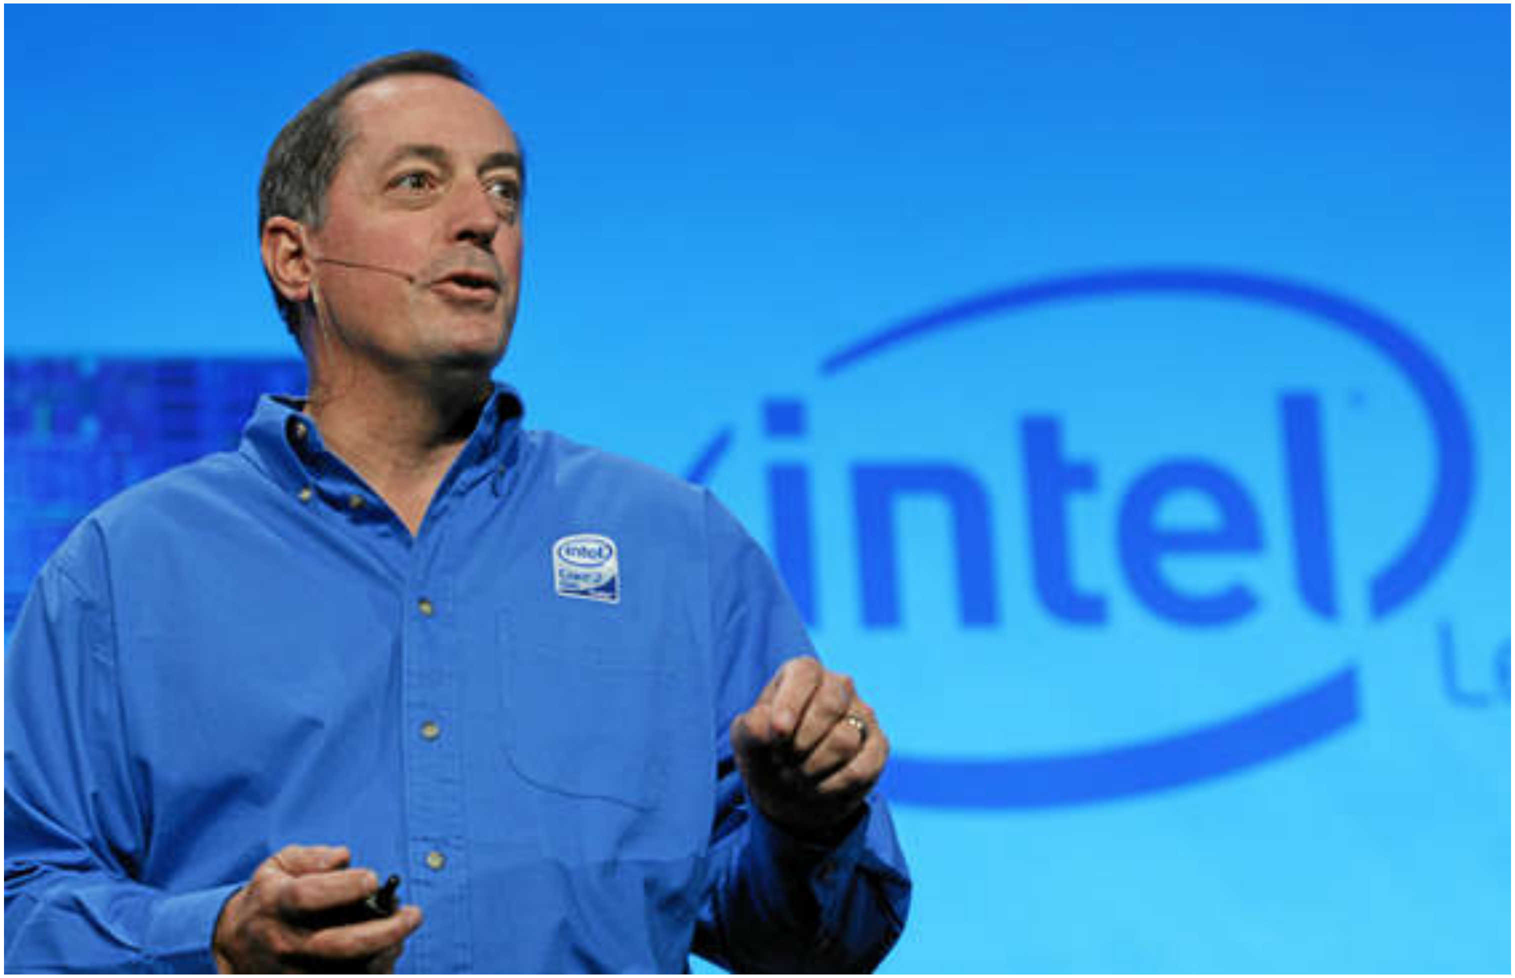

What were our aims?

Abstract of LUG 2011 Abstract

…create consistent repeatable test processes…

…maximize the value of the distributed community testing required by the large test matrix.

…show how community testing can be consistent enough to allow for process improvement whilst being flexible enough to enable each partner to operate independently.

…component based approach allows considerable flexibility in the test environment whilst providing the closed loop test process required by any high quality software product.

…show how test results can become part of the landing collateral for a patch…

What have we done?

We’ve done quite a lot…

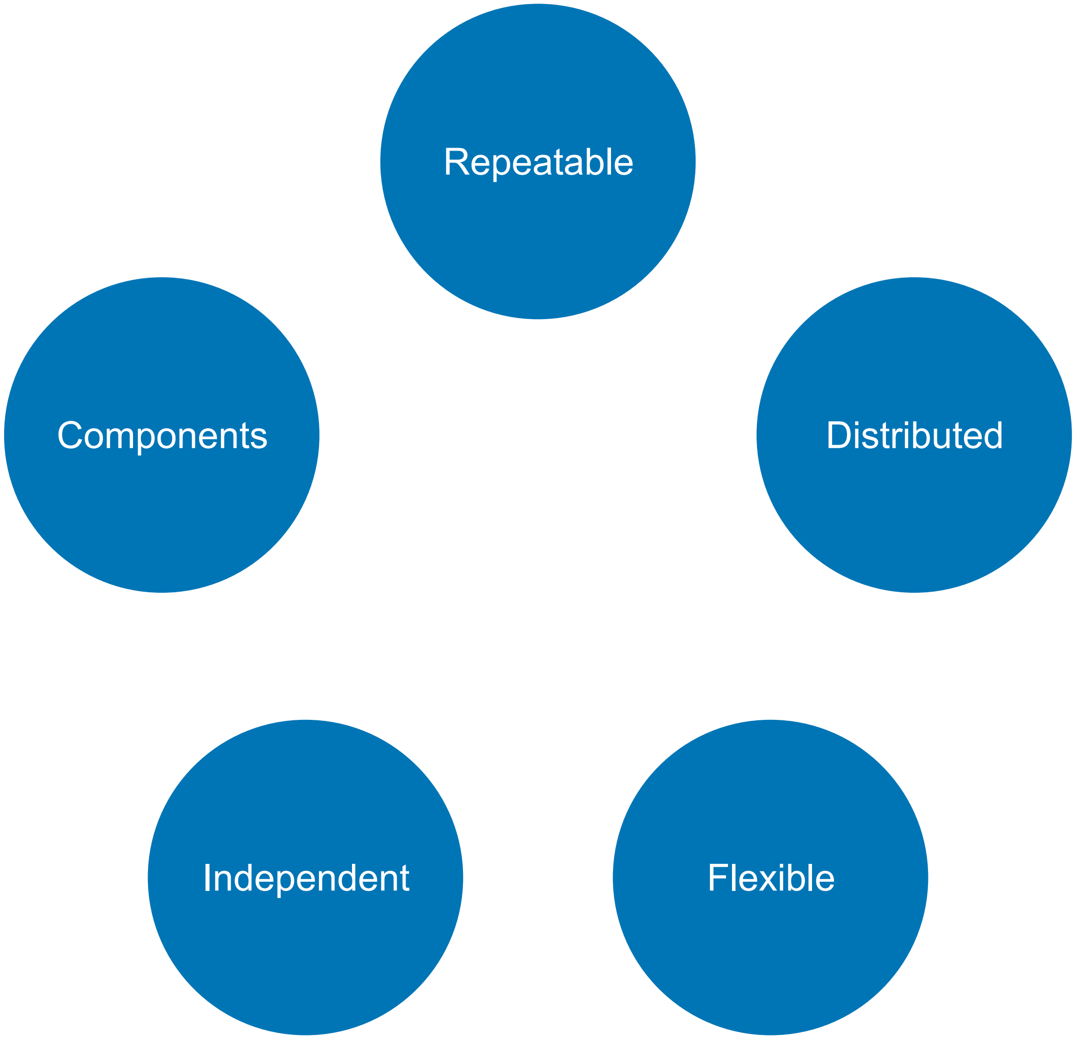

Repeatable

Components

Distributed

Independent

Flexible

Distributed

- Co-operative, Scalable Cloud
  - Co-operative test clusters in Colorado, DuPont, New Mexico, Jülich and Indiana University (OpenSFS)
  - Clusters share the work or work independently
  - Clusters provide different attributes
  - Additional scale test at LLNL

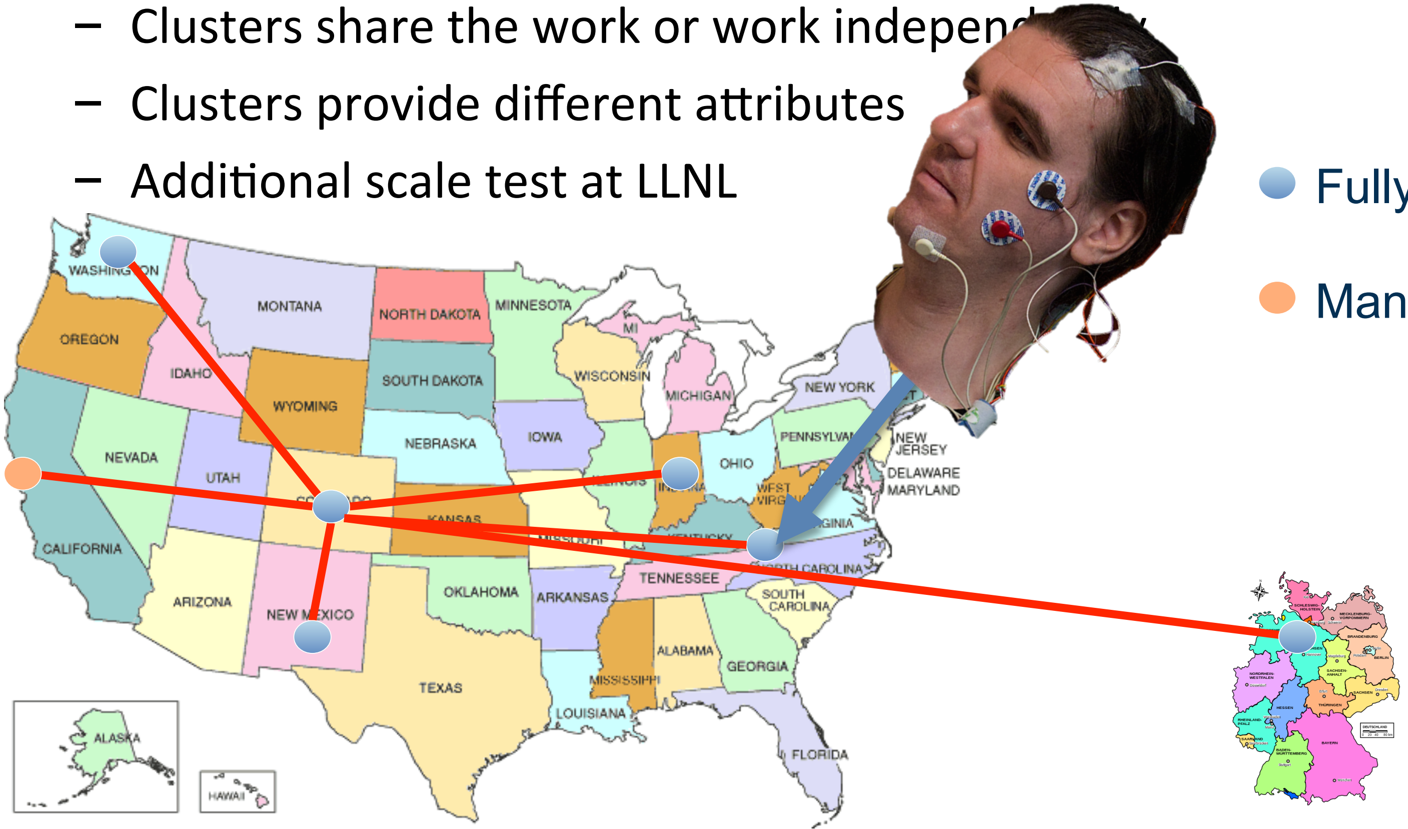

● Fully Automatic

● Manual Test

Flexible

- Test systems can be anything that supports a Lustre* File System
  - Test requirements will guide the test to the correct locations
    - OpenSFS NRE work will utilize the Indiana Cluster
    - Large LUN will be tested in Jülich
    - Larger Clusters will be tested in Du Point or New Mexico
  - A single build can be tested at multiple sites
    - Master reviews tested for DNE on OpenSFS
- Flexible Test Capability
  - Lustre e2fsprogs-reviews against a Lustre file system
  - Lustre Manual tested by building
  - HSM will be tested on HSM capable clusters
- Every change should have some verification process

Consistent aka Repeatable

- Developers agnostic to the location of test
  - Tests may occur at many sites with equal results
  - Developers don't care – although they do know
- Every test session occurs on a new purpose built cluster
  - Complete ground up install of OS, Lustre and file systems
  - If bad state detected then the cluster is completely rebuilt
  - Every test records the version of all system components
  - Any historic test system could be rebuilt

Components

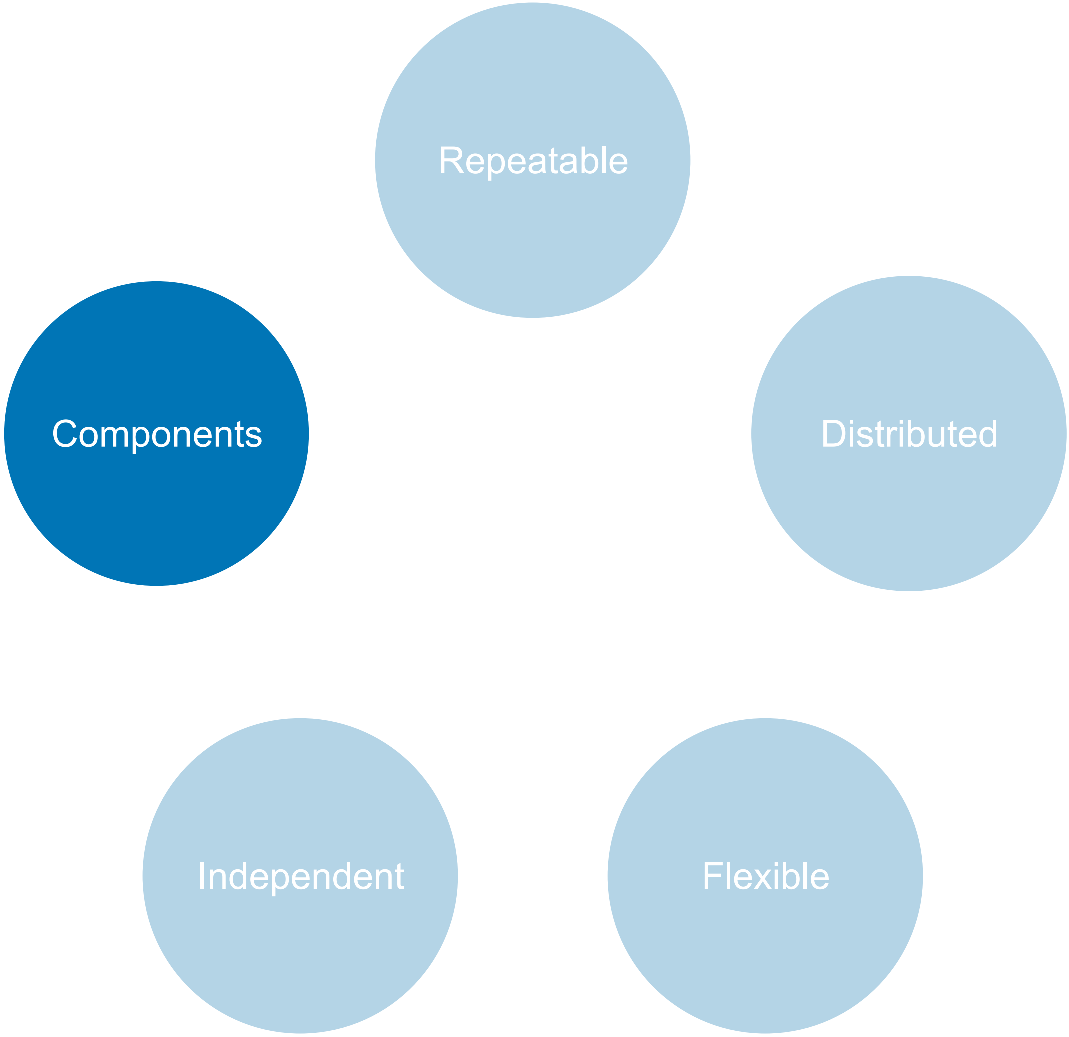
Repeatable
Components
Distributed
Independent
Flexible

Components

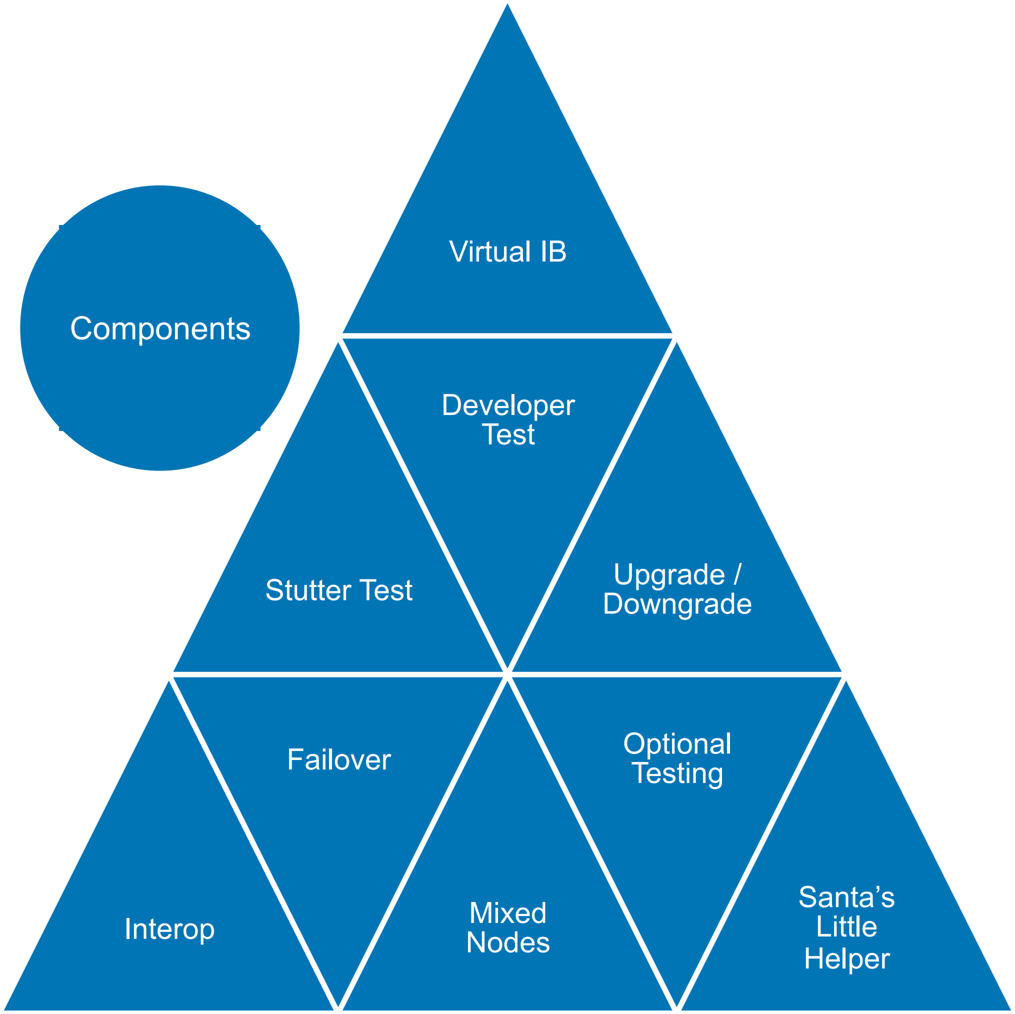
Components

Virtual IB

Developer Test

Stutter Test

Upgrade / Downgrade

Failover

Optional Testing

Interop

Mixed Nodes

Santa’s Little Helper

Virtual Infiniband

- Frank Heckes working closely with Mellanox to rollout virtual IB
  - Still using RC Mellanox code but facility rolling out
- Cheaper larger clusters mixing IB and vIB nodes
  - Not a replacement for real IB but a very good analogy
- Most testing now possible on IB rather than Ethernet

Stutter Testing

- Modern CPU's have a clock-range
  - Different parts of the cluster will be clocking at different speeds
  - Even when the system is running 'flat-out'
  - Oleg realized we could simulate this very well by using many more virtual-Cores than real Cores
- Our tests now allow us to vary the effect between test sessions
  - Range varies from repeatable to all over the place…

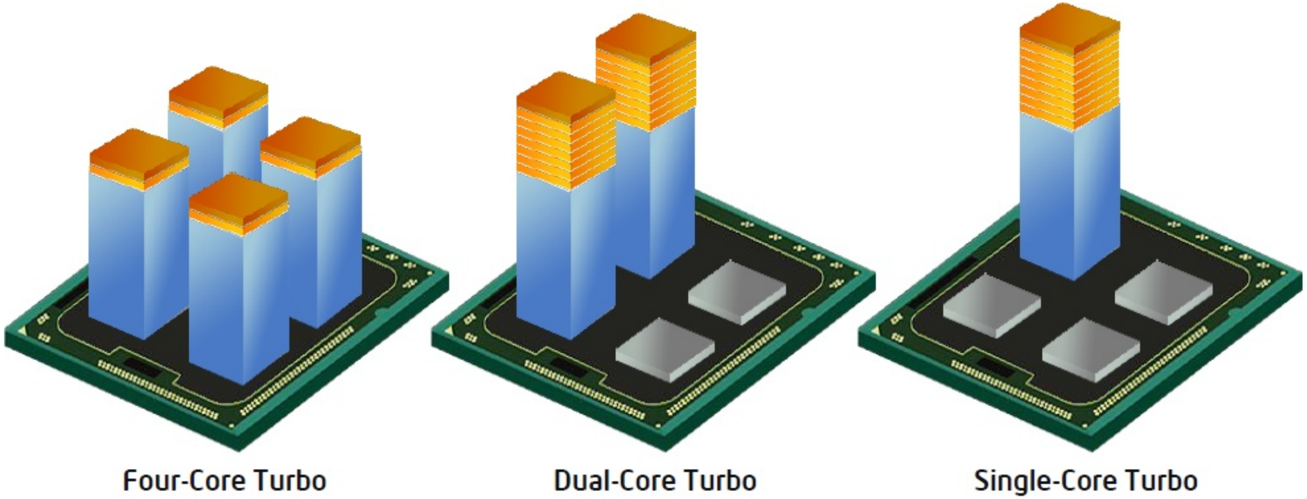

Upgrade / Downgrade / Interop Testing

- Test any compatible version with any compatible version
  - Standard automated testing covers the whole Matrix
- Complete upgrade path testing matrix
  - From any version to any version
  - Up and Down
- All features can be tested as part of interop and up/down grade testing

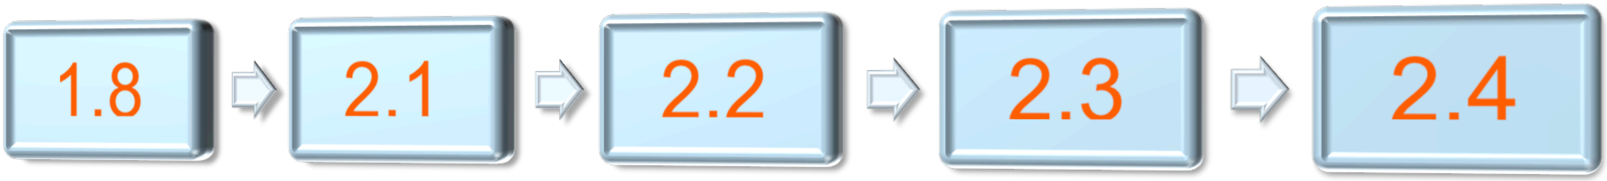

Developer Test

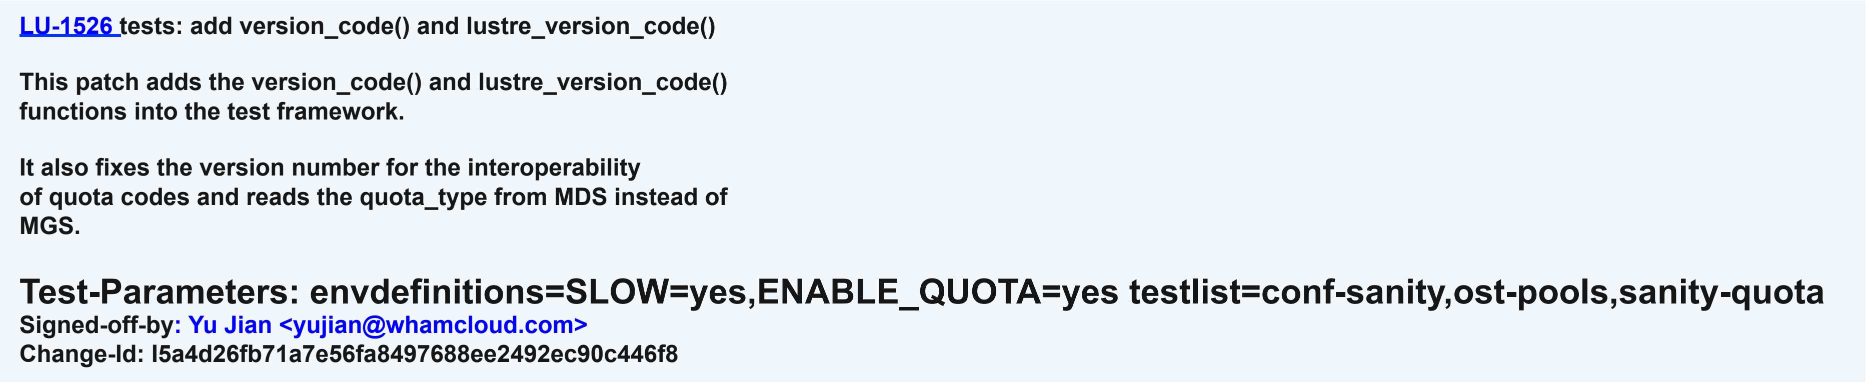

LU-1526 tests: add version_code() and lustre_version_code()

This patch adds the version_code() and lustre_version_code()
functions into the test framework.

It also fixes the version number for the interoperability
of quota codes and reads the quota_type from MDS instead of
MGS.

Test-Parameters: envdefinitions=SLOW=yes,ENABLE_QUOTA=yes testlist=conf-sanity,ost-pools,sanity-quota
Signed-off-by: Yu Jian <yujian@whamcloud.com>
Change-Id: I5a4d26fb71a7e56fa8497688ee2492ec90c446f8

- Developer test actually running
  - 286 developer test runs
    - 1 per day over last year
    - 6 per day over last month
- Regular debug method
- Target regular development method

Flexible Optional Testing

- Minimum `bar` + as much as possible
  - Always test to the bar
    Then test as much as possible
- Some additional tests are carried out if capacity allows
  - Required standard always achieved
  - Higher standard enabled

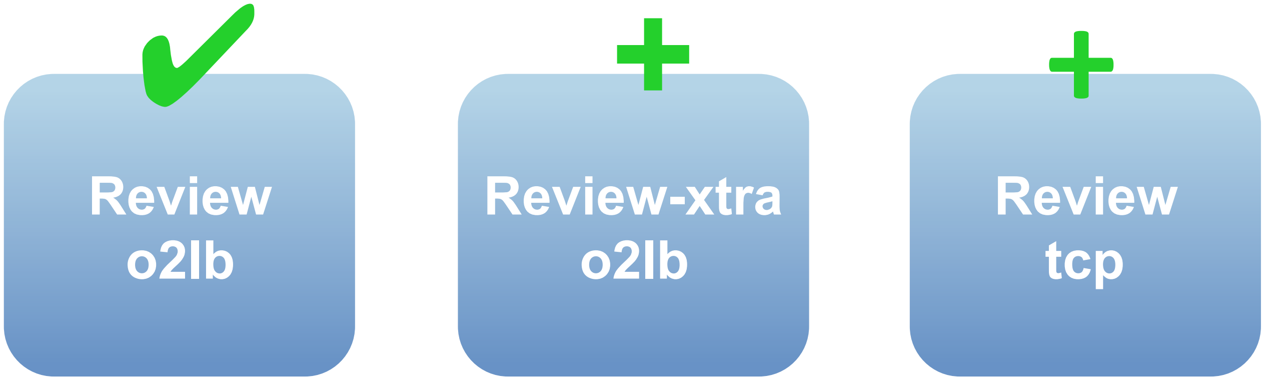

✓ Review o2lb

+ Review-xtra o2lb

+ Review tcp

Where are we going?

Verify behavior as well as validating function

- Most Lustre testing is functional
  - 90% of the test scripts test a particular function
- We can sometimes loose sight of the behaviour we want
- More focus on verifying we have the behaviour we want
  - Develop tool sets and methods to continuously verify behaviour
    - Performance
    - Scalability
    - Reliability
    - …

intel®