

# OFED Status & Plans Sonoma 2009



OPENFABRICS  
ALLIANCE

Betsy Zeller, Tziporet Koren

# Agenda

- Linux OFED components
- Releases done in last year:
  - OFED 1.3.1
  - OFED 1.4
- 2009 Plans:
  - OFED 1.4.1
  - OFED 1.5
- How to contribute

# Linux OFED Components

## OFA Development

- HCA/NIC Drivers
  - IB: IBM, Mellanox, QLogic
  - iWARP: Chelsio, Intel
- Core: Verbs, mad, SMA, CM, CMA
- IPoIB
- SDP
- SRP and SRP Target
- iSER and iSER Target
- RDS
- NFS-RDMA
- Qlogic\_VNIC
- uDAPL
- OSM
- Diagnostic tools

## Add on

- Bonding module
- Open iSCSI
- MPI Components
  - MVAPICH
  - Open MPI
  - MVAPICH2
  - Benchmark tests

## Tested with

- Proprietary MPIs: Intel, HP
- Proprietary SMs: Cisco, Voltaire, Qlogic

# Agenda

- Linux OFED components
- Releases done in last year:
  - OFED 1.3.1
  - OFED 1.4
- 2009 Plans:
  - OFED 1.4.1
  - OFED 1.5
- How to contribute

# OFED 1.3.1

- OFED 1.3.1 release on June 3, 2008
  - Added support for RedHat EL 5.2 and SLES 10 SP2
  - Fixed several critical bugs
- Distro integration:
  - Red Hat AS 4.7 and RHEL 5.2, SLES10 SP2
- Used in Intel ® Cluster Ready Solutions
- Passed Oracle 11g certification with RDS

# OFED 1.4

## ➤ General Info

- Released in December 10, 2008
- Passed in the Interoperability event in Nov 2008
- Added support for CentOS and OEL (Oracle Enterprise Linux)
- Kernel base 2.6.27

## ➤ Distro integration:

- SLES 11
- RHEL 4.8, 5.4 (not released yet)

## ➤ Used in Intel ® Cluster Ready Solutions

# OFED 1.4 Features

- New: BMME verbs (fast memory thru send queue (FRWR); Local invalidate send work requests; Read with invalidate)
- New: iSer Target
- New: NFS-RDMA – as technology preview
- New: VPI support – Eth and IB for ConnectX

# OFED 1.4 Features – Cont.

- IPoIB:
  - LRO and LSO for Datagram mode
  - Improved Bonding failover response time
- uDAPL:
  - Socket CM for scalability and interop with Windows
  - UD extensions
- Qlogic\_vnic:
  - Hot swap of EVIC and dynamic update of existing connections with QLogic dynamic update daemon.
  - Performance improvements of Ethernet broadcast & multicast traffic.



# OFED 1.4 Features - Cont.

- New management package (ver 3.2):
  - OpenSM
    - Cashed routing
    - Multi lid routing balancing for updn/minhop routing algorithms
    - Preserve base lid routes when LMC > 0
    - OpenSM configuration unification
    - IPv6 Solicited Node Multicast addresses consolidation
    - Routing Chaining
    - Failover/Handover improvements: Query remote SMs during light sweep
    - Ordered routing paths balancing
  - ibutils:
    - Report created in CSV format
    - Congestion Control in ibutils
  - Diagnostic tools: ibnetdiscover library - to accelerate another tools

# OFED 1.4 Features - Cont.

- New MPI versions:
  - OSU MVAPICH 1.1.0
  - Open MPI 1.2.8
  - OSU MVAPICH2 1.2p1
  - Tests: Updated IMB 3.1

# OFED 1.4 OS Matrix

- kernel.org: **kernel 2.6.26 and 2.6.27**
- Novell
  - SLES 10
  - SLES 10 SP1 (up1)
  - SLES 10 SP2
- Redhat
  - RHEL 4 (up4, up5, up6, **up7**)
  - RHEL 5 (up1, up2)
- OEL
  - **OEL 5**
- Free distros (with limited QA):
  - Open SuSE 10.3
  - **Fedora Core 9**
  - Ubuntu 6.06 (with RPM package installed)
  - **CentOS 5.2**

\* *new for OFED 1.4 in bold*

# Agenda

- Linux OFED components
- Releases done in last year:
  - OFED 1.3.1
  - OFED 1.4
- 2009 Plans:
  - OFED 1.4.1
  - OFED 1.5
- How to contribute

# OFED 1.4.1- Main features

- Support for RHEL 5.3 and SLES 11
- NFS/RDMA in beta
  - OSes: RHEL 5.2, 5.3 and SLES 10 SP2
- Open MPI 1.3.1
- RDS with iWARP support in beta
- VPI ConnectX IB/Eth port sensing
- Critical bug fixes

# OFED 1.4.1

## ➤ Schedule:

- RC1 - Mar 4 - done
- RC2 - Mar 19 - done
- RC3 - Apr 2
- GA - Apr 20

# OFED 1.5 Features Plans

- Kernel base: 2.6.30
- Add support for RedHat EL 5.4 and EL 4.8
- Kernel verbs: Multiple Event Queues to support Multi-core CPUs
- NFS/RDMA – GA
- RDS from the kernel; support for iWarp – GA
- SDP – Performance improvements: small and medium messages BW, reduced jitter, GA quality
- Support for Mellanox vNIC (EoIB) and FCoIB with BridgeX device
- New MPI features – details in MPI session
- More features according to requirements that will be raised here ...

# OFED 1.5 – Management Features



- Unify API with Windows
- OSM:
  - Fat-tree enhancements:
    - Connect roots
    - Credit loop-free multicast routing with managed switches
  - SM handover – enable SM on every node
  - Shadow SA DB
  - M\_Key management
- More details in OpenSM Update



# OFED 1.5 OS Matrix

- kernel.org: **kernel 2.6.29 and 2.6.30**
- Novell
  - SLES 10
  - SLES 10 SP1 (up1)
  - SLES 10 SP2
  - SLES 11
- Redhat
  - RHEL 4 (up4, up5, up6, up7, **u8**)
  - RHEL5 no updates, up1
  - RHEL 5 (up2, u3, **up4**),
- OEL
  - OEL 5
- Free distros (with limited QA):
  - Open SuSE 10.3
  - Fedora Core 9
  - Ubuntu 7 (with RPM package installed)
  - CentOS 5.2, 5.3

- *new for OFED 1.5 in **bold***
- *drop support for items in **blue***

# OFED 1.5 Schedule

## ➤ Preliminary Schedule

- Development tree opened when 2.6.30-rc1 is available
  - People can start development now against 2.6.29 Linux kernel
- Feature Freeze: May 7, 09
- Alpha Release: May 12, 09
- Beta Release: Jun 9, 09
- RC1: Jun 25, 09
- RC2-RCx: About every 2 weeks as needed
  - We usually have ~6 RCs
- Release: Sep 15, 09

# What is an RC?

- RC = Release candidate – something pretty close to what we'd like to release.
- An early RC will be sent for interoperability testing.
- Not the time to complete your new feature!
- This is the opportunity for testing and fixing bugs.

# How to contribute?

- Developing new code and features
- Bug fixes
- Performance tuning
- Contribute backports for new OSes
- Doing QA and testing
- Sending patches and comments to the mailing lists:
  - [ewg@lists.openfabrics.org](mailto:ewg@lists.openfabrics.org) – **OFED specific only**
  - [general@lists.openfabrics.org](mailto:general@lists.openfabrics.org) – **General development**
- Opening bugs in Bugzilla (<https://bugs.openfabrics.org/>)
  - When opening a new bug you should choose OpenFabrics Linux
  - Old bugs must be tested with new releases and updated on bugzilla
- Participate in EWG bi-weekly meetings
  - Meeting minutes on the web:  
[http://www.openfabrics.org/txt/documentation/linux/EWG\\_meeting\\_minutes/](http://www.openfabrics.org/txt/documentation/linux/EWG_meeting_minutes/)