

OpenSM and InfiniBand Management Update



OPENFABRICS
ALLIANCE

Sasha Khapyorsky

Agenda

- General
- New things
 - OFED 1.4
 - The current
- Plans
 - OFED 1.5
 - Long terms
- Wishes
- Thanks

➤ Git tree

- <git://git.openfabrics.org/~sashak/management>

➤ Downloads

- <http://www.openfabrics.org/downloads/management/>

➤ Daily builds

- <http://www.openfabrics.org/downloads/management/daily>

- OFED 1.4.x versions:
 - libibcommon-1.1.2, libibumad-1.2.3, libibmad-1.2.3, opensm-3.2.6, infiniband-diags-1.4.4, ibsim-0.5+
- The current versions:
 - libibumad-1.3.1, libibmad-1.3.1, opensm-3.3.1, infiniband-diags-1.5.1
- Release Notes:
 - `opensm/doc/opensm_release_notes-3.X.txt`

IB Management core libraries

➤ libibcommon

- is gone – merged over libibumad and libibmad

➤ libibumad

- Sysfs routines are here now (sysfs.c)
- Couple of fixes

➤ Libibmad

- Couple of fixes
- Almost thread safe now – old API is deprecated
- WinOF porting
- Couple of fixes and MAD encoding enhancements

InfiniBand Diagnostic tools

➤ All tools

- Unified command line interface
- Converted to use new thread safe libibmad API
- WinOF porting
- Many fixes and improvements

➤ Perfquery

- XmitWait counter support
- PortXmitDataSL, PortRecvDataSL counters support
- Ports looping option (--loop_ports)

InfiniBand Diagnostic tools

- saquery
 - Converted to use libibumad
 - SM_Key support
 - More queries (PIR, LFTR, MFTR)
 - Query parameters (--mtu, --sgid, etc.)
- vendstat
 - More vendor specific stuff - counter group info
- ibsysstat
 - Use RMPP – large data buffers are supported now
- new tool – ibsendtrap

➤ General

- Many fixes, optimizations, improvements and cleanups

➤ User Interface

- Unified configuration file
- Configuration reloading on the fly
- Improved Plugin interface – multiple plugins are supported

➤ Multicast

- Ipv6 Solicited Node Multicast addresses consolidation

OpenSM (Routing)

➤ Routing

- Cached routing
- Routing chaining - `-R ftree,updn,minhop`
- LMC improvements:
 - Preserve base LIDs routes
 - Ensure LMC paths balancing over different switches/chassis
- Ordered paths balancing
 - Ports are sorted by switch loads
 - Port order file option (`--guid_routing_order_file` option)
- Mesh geometry analysis for LASH

OpenSM (Routing)

- Routing (cont.)
 - Paths balancing over multiple links for LASH
- Low Level routing tuning
 - Port IDs for Up/Down
 - Min hop weights
- Hacks
 - Connecting root nodes with Up/Down
 - Connecting IO nodes with FastTree

➤ Handover improvements

- OpenSM will query Standby SMs periodically
- Standby OpenSM notifies Master SM about priority change (Trap 144)

➤ SA Scalability

- Shadow SA DB
- More optimal SM/SA DB update interaction
- Other potential SA optimizations
 - PR/MCMR for homogeneous network
- SA DB replication

➤ OpenSM User Interface

- OpenSM configuration/comands framework
 - configuration file, command line options, console, etc..
- Secure console

Plans

- Incremental routing
- MGIDs compression
 - An option to map many MGIDs to single MLID
- MKey management
- QoS manager integrations
 - Share port group definition with Partition Manager
 - Low level QoS parameters: SL2VL, VLArb, etc
 - QoS/LASH co-existence
 - SL2VL setup optimization

Plans

- **Ibnetdiscover library**
 - Should speedup other diag tools dramatically
- **OpenSM diagnostic capabilities**
- **WinOF porting**

Plans: Long Term

- Congestion management
- IB Router
- Distributed SA

Wishes

- Even more collaboration is needed
 - More feedback, testing, development are desired
- Thoughts, Ideas?
 - Not always we need to wait for next Sonoma
 - general@lists.openfabrics.org
 - sashak@voltaire.com

Thanks

➤ `git log ofed_1_3..master --pretty=format:%an | sort -u`

Albert L. Chu

Arlin Davis

Dale Purdy

David A. McMillen

Doron Shoham

Dotan Barak

Doug Ledford

Eli Dorfman

Hal Rosenstock

Ira K. Weiny

Itai Baz

Jack Morgenstein

Jeremy Brown

Keshetti Mahesh

Line Holen

Matthias Blankenhaus

Max Matveev

Mike Heinz

Nicolas Morey

Chaisemartin

Or Gerlitz

Ralph Campbell

Robert Pearson

Ron Livne

Sasha Khapyorsky

Sean Hefty

Stan Smith

Timothy A. Meier

Vincent Ficet

Yevgeny Kliteynik

Thank You