



IBM Linux Technology Center

September 7, 2006

# Infiniband on Xen Update



Nivedita Singhvi

# Virtualizing Infiniband

---

- Work done by Jiuxing Liu (IBM Research) and
  - ▶ Bulent Abali (IBM Research), Wei Huang, D.K. Panda (Ohio State)
- Virtualized infiniband driver for an infiniband HCA adapter
- Initial work started around 3.0.0 timeframe
- Presented at January, 2006 Xen Summit by Jiuxing
- Opened up source code on mercurial tree:
  - ▶ <http://xenbits.xensource.com/ext/smart-io.hg>
- IB Xen Wiki page:
  - ▶ <http://wiki.xensource.com/xenwiki/XenSmartIO>



# Virtualizing Infiniband - Tasklist

---

- Implement Xenbus/xenstore support
  - ▶ Eliminate need for out-of-band tcp communication
  - ▶ manually set things up at module load time
- Fix code for cleanup/shutdown/error handling
  - ▶ Lots of problems here
- Improve xen interdomain communication support in data\_ch.[hc]
- Fix granting access of UAR pages to user domains
  - ▶ Currently this is done by disabling permission checks in Xen
- Fix security checks/holes
- Fix portability problems ( currently only x86)



## Virtualizing Infiniband – Tasklist (2)

---

- Reorganize the code to isolate hardware HCA dependent part
  - ▶ support other HCAs
- Multiple physical/virtual HCAs support
  - ▶ Currently we support only 1 physical HCA and present only one virtual HCA to each domain
- Currently only Mellanox MT23108 compatible cards are supported
- IB management support
- Checkpointing/Migration support



# Virtualizing Infiniband Update

---

- Need to leverage community to help continue work
  - ▶ *Lot of interest after Jan summit*
  - ▶ *Not yet getting top priority / enough cycles*

