

ATOM Server Setup

August 6, 2014

Satoshi Matsushita , Y. Hasebe

Stanford Univ. / NEC

Translated from NEC proposal rev6 by Hasebe



Table of Contents

1. Sever architecture overview	p. 3 - 6
2. NIC assignment for RAMCloud on ATOM server	p. 7
3. Network connection of ATOM cluster	p. 8
4. ATOM Server Setup	
4.1. Network Connection in a Chassis	p. 9
4.2. VLAN Reconfiguration	p.10-13
5.	



NEC Micro Modular Server


- Globally announced on May 20, 2014: (Press release : [NEC raises the bar for high density IT solution platforms for the public and private cloud](#))

Chassis: Redundancy (power supply, Networks, Fans) + Hot Swappable

- 2U in standard 19inch rack
- Up to 46 **Atom** server with
32GB DRAM / 128GB SSD / 2x 2.5GbE
- 2x 230Gbps **switch** (FM5224), 4x 40Gbps uplinks
- Chassis Total: 1.4TB DRAM, 5.8TB SSD
- max. 2kW

16 chassis in a rack:


- 736 ATOM Servers : 5.9k core, 23TB DRAM, 92TB SSD:
50TOp/s, 20TFlops, DRAM 368TB/s, SSD 647GB/s



ATOM Server Module

A Server Module

Block Layout

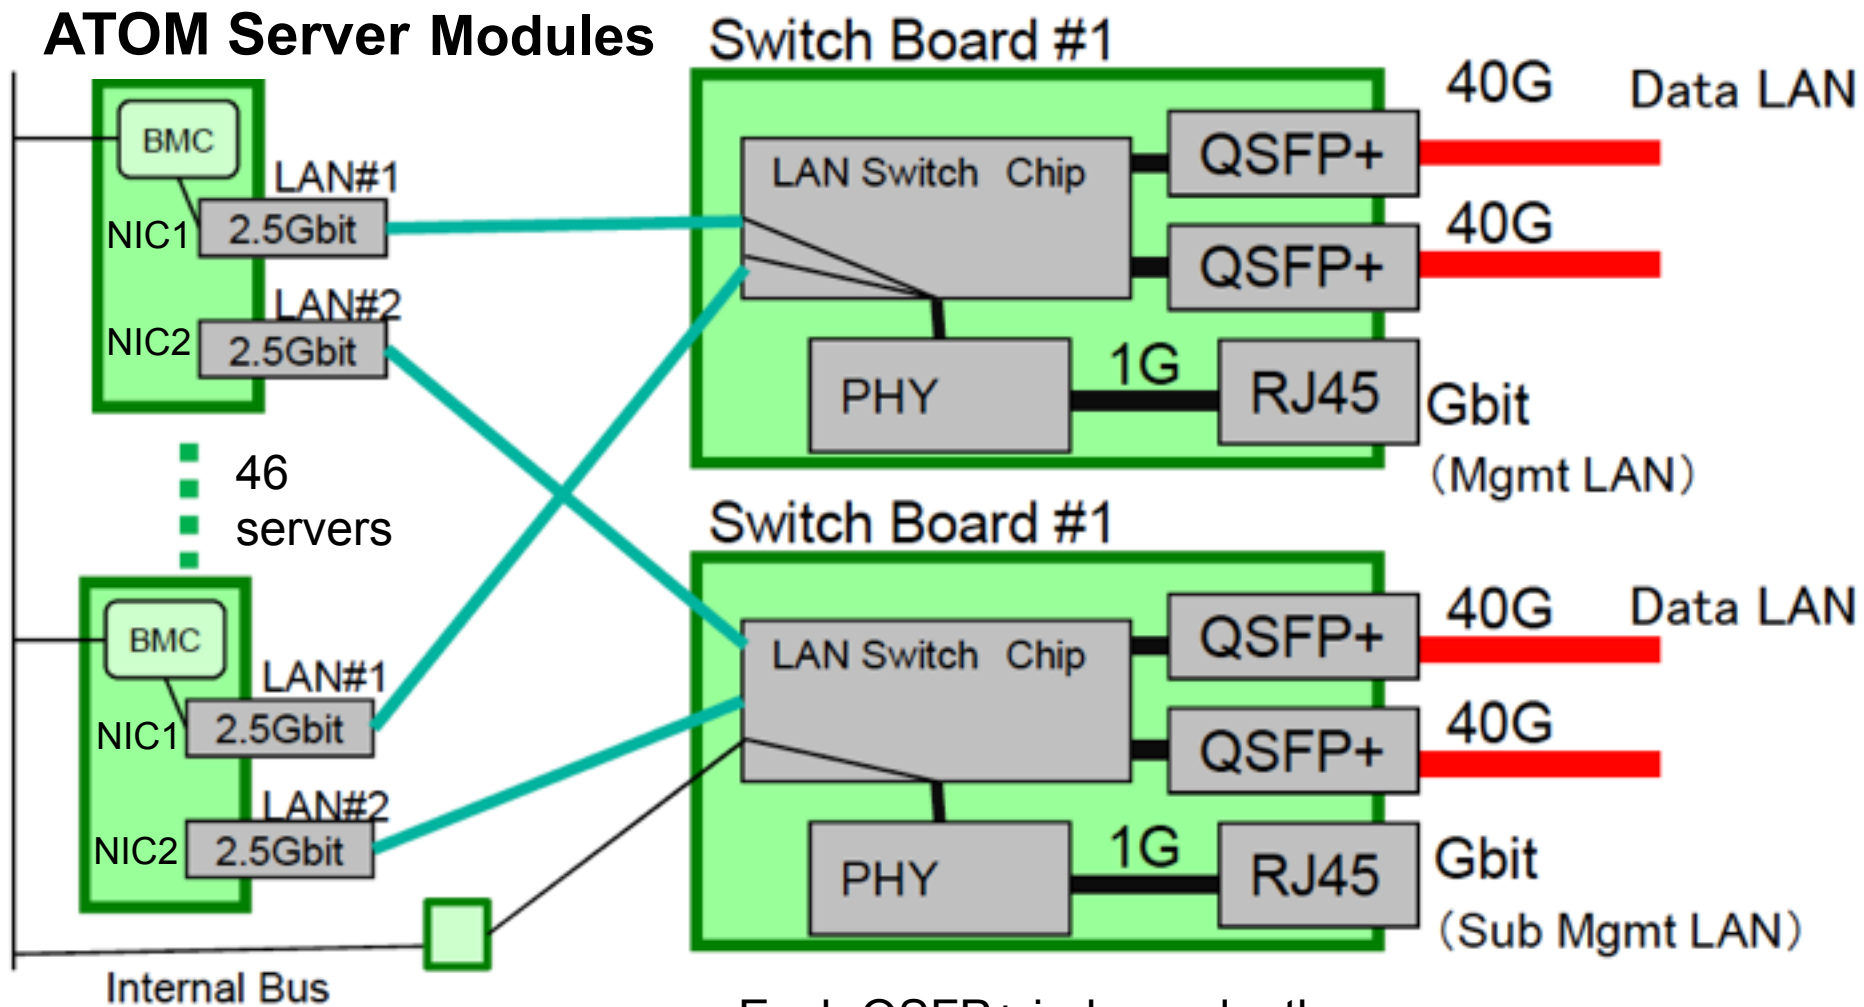


【SPECIFICATIONS】

- 1x CPU(Atom™ C27xx)
- 4x SO-DIMM (Max 32GB) w. ECC
- 1x mSATA SSD (128GB)
- 1x BMC
- 1x SATA3 (To mSATA SSD)
- 2x SATA2 (To storage board)
- 2x 2.5Gbit LAN

Processor	Cores	Frequency	Power
C2750	8C / 8T	2.4GHz	20W
C2730	8C / 8T	1.7GHz	12W

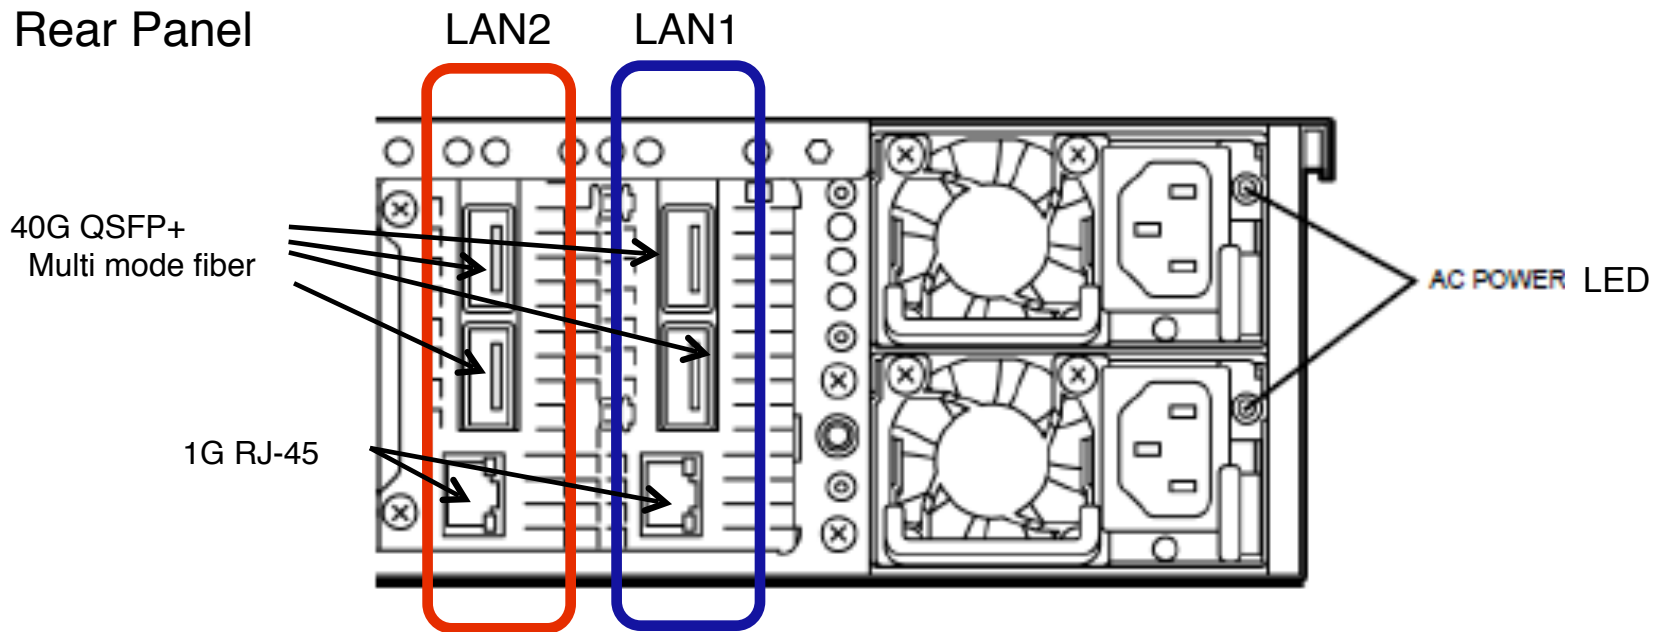
Connection in a Chassis



Each QSFP+ independently configurable as ether 1x 40G or 4x 10G

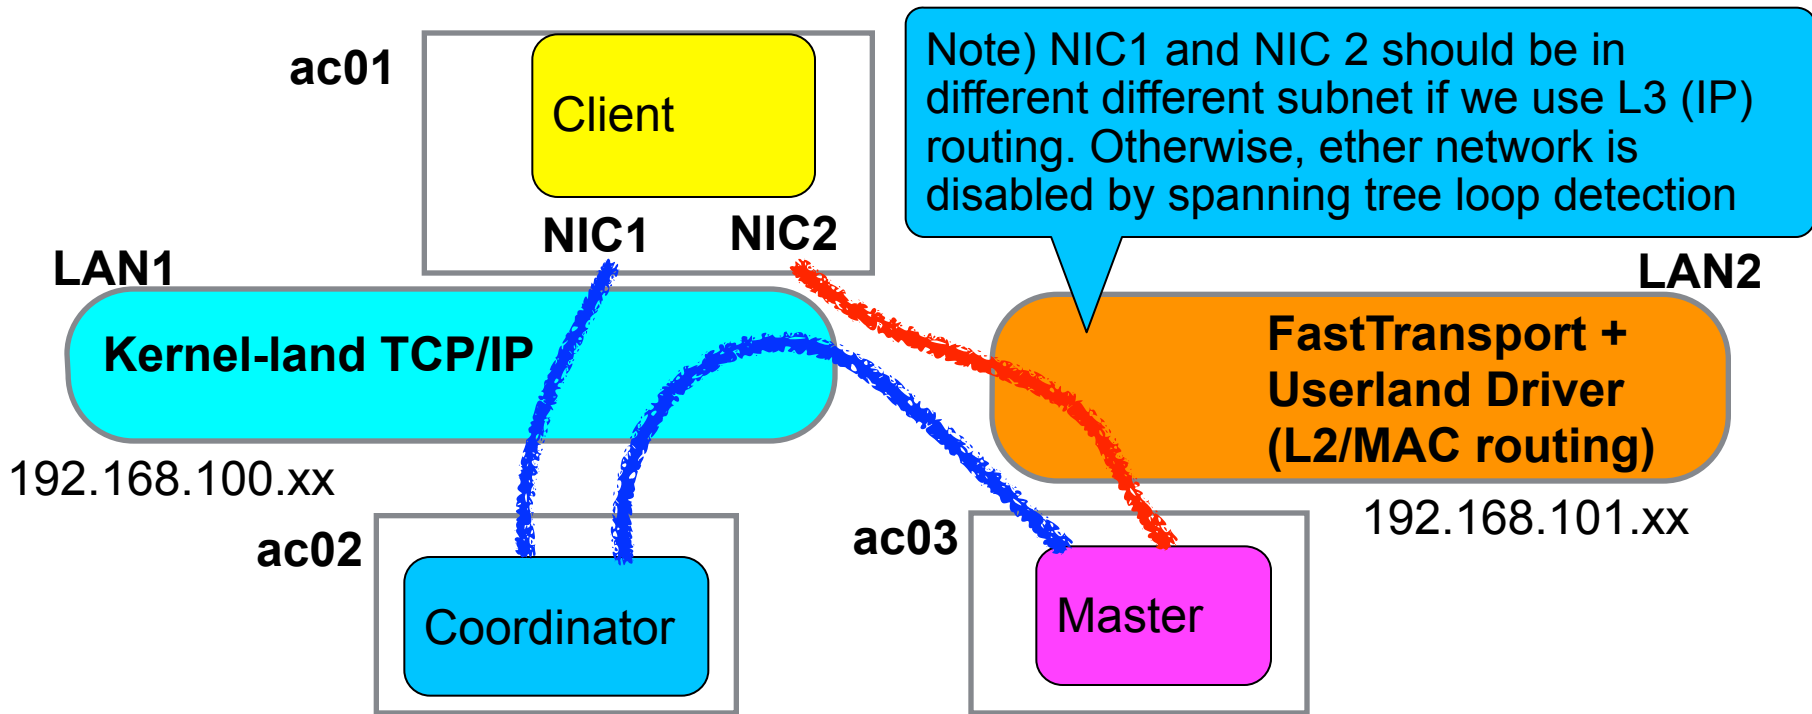
Network Ports

Rear Panel



RAMCloud with Userland Driver

- Using userland driver only for critical Master-Client data path.
- Developed userland driver for NIC2 with L2 (MAC) routing using Intel DPDK.
 - No Need to assign IP address to LAN2 with DHCP.



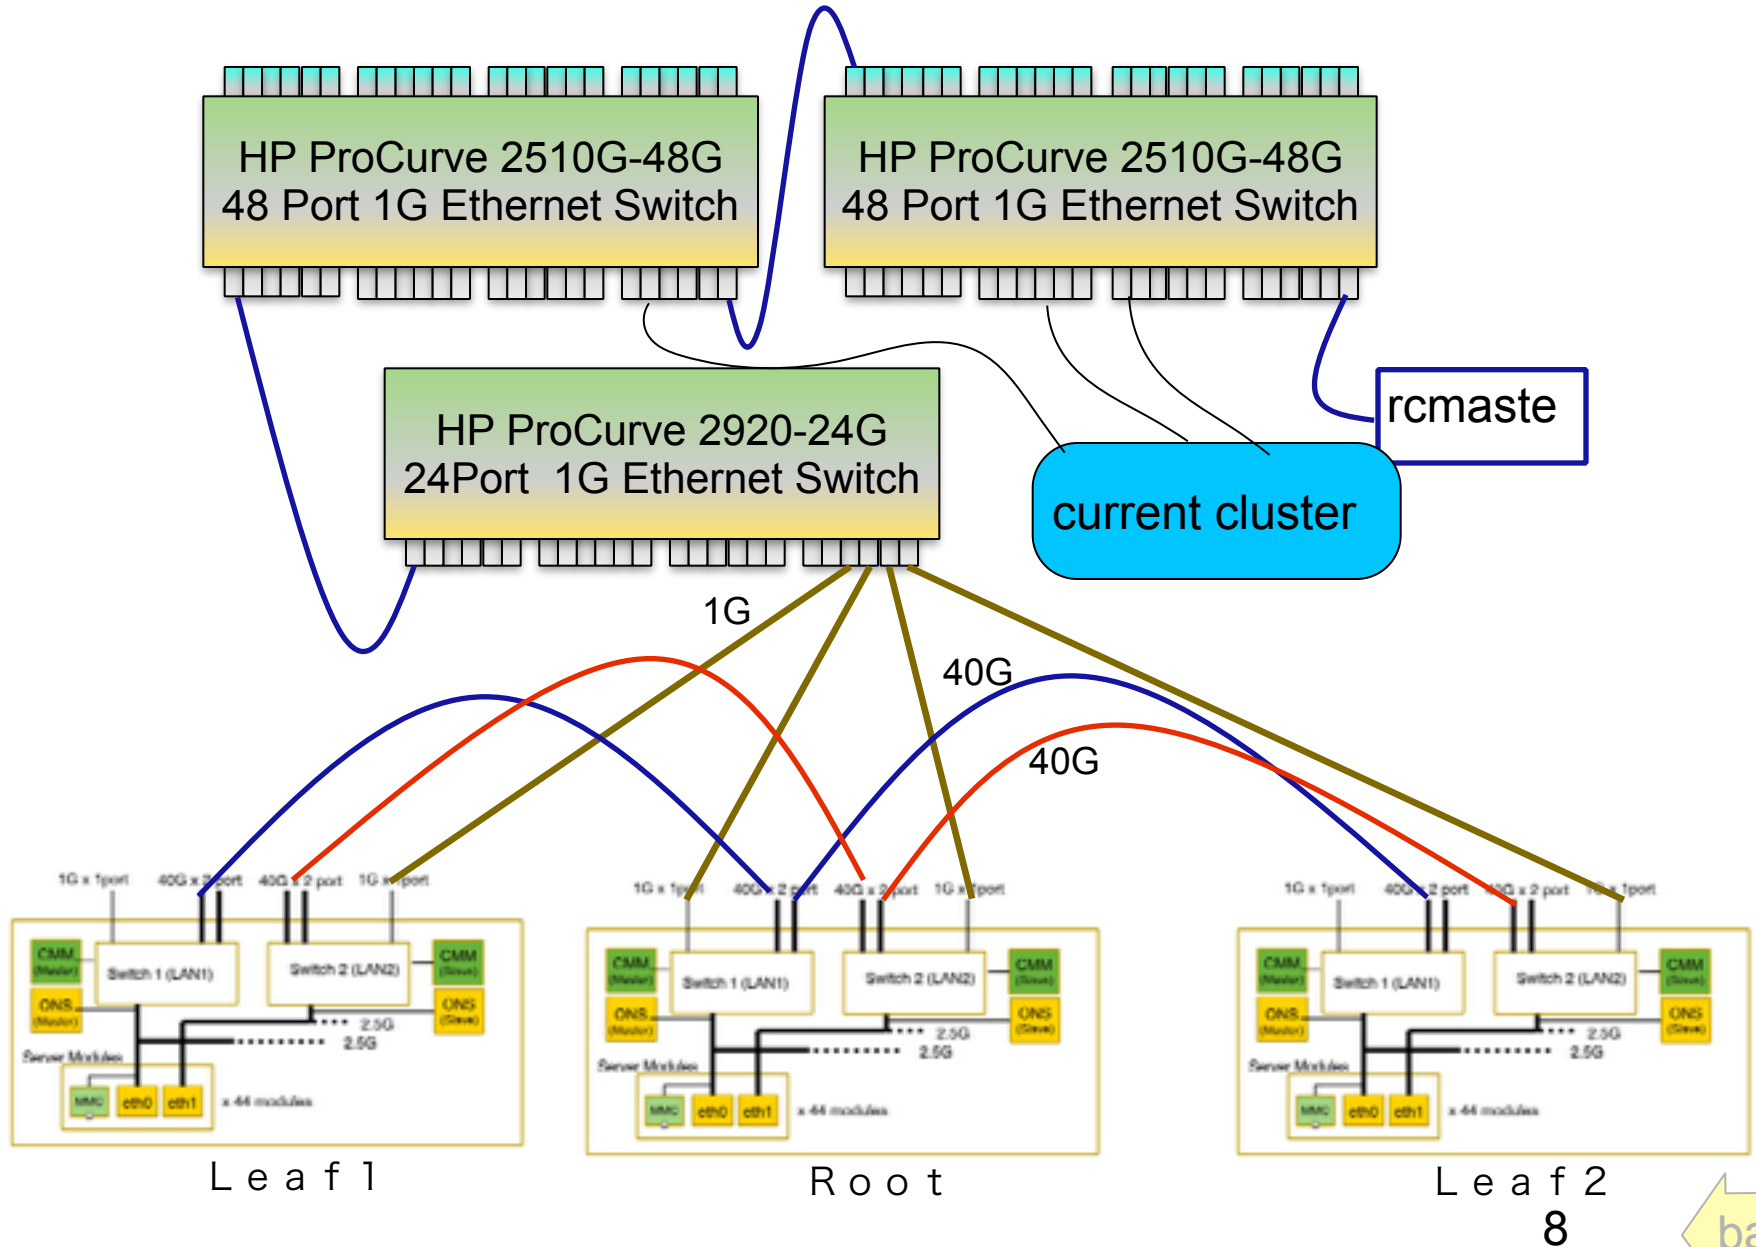
Command Line:

```
$ coordinator -C tcp:host=192.168.100.31,port=12246
```

```
$ server -C tcp:host=192.168.100.31,port=12246 -L fast+dpdk:host=192.168.101.29,mac=94:DE:80:AB:01:79 -r 0
```

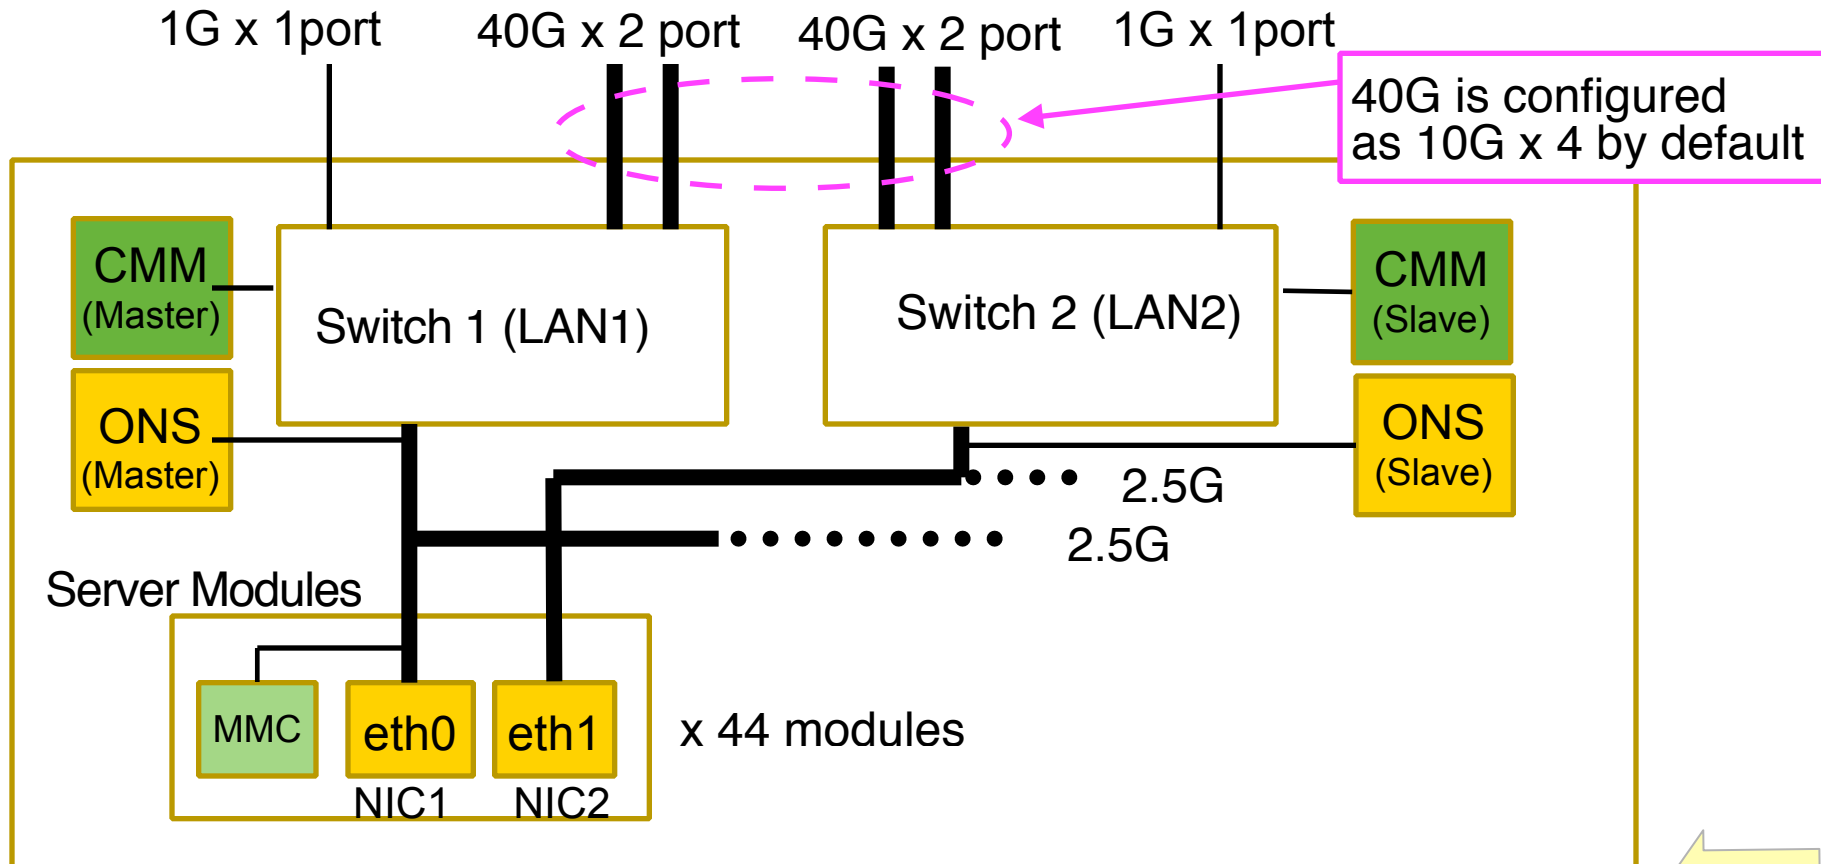
```
$ ClusterPerf -C tcp:host=192.168.100.31,port=12246 --numClients 1 basic
```

ATOM Cluster Network



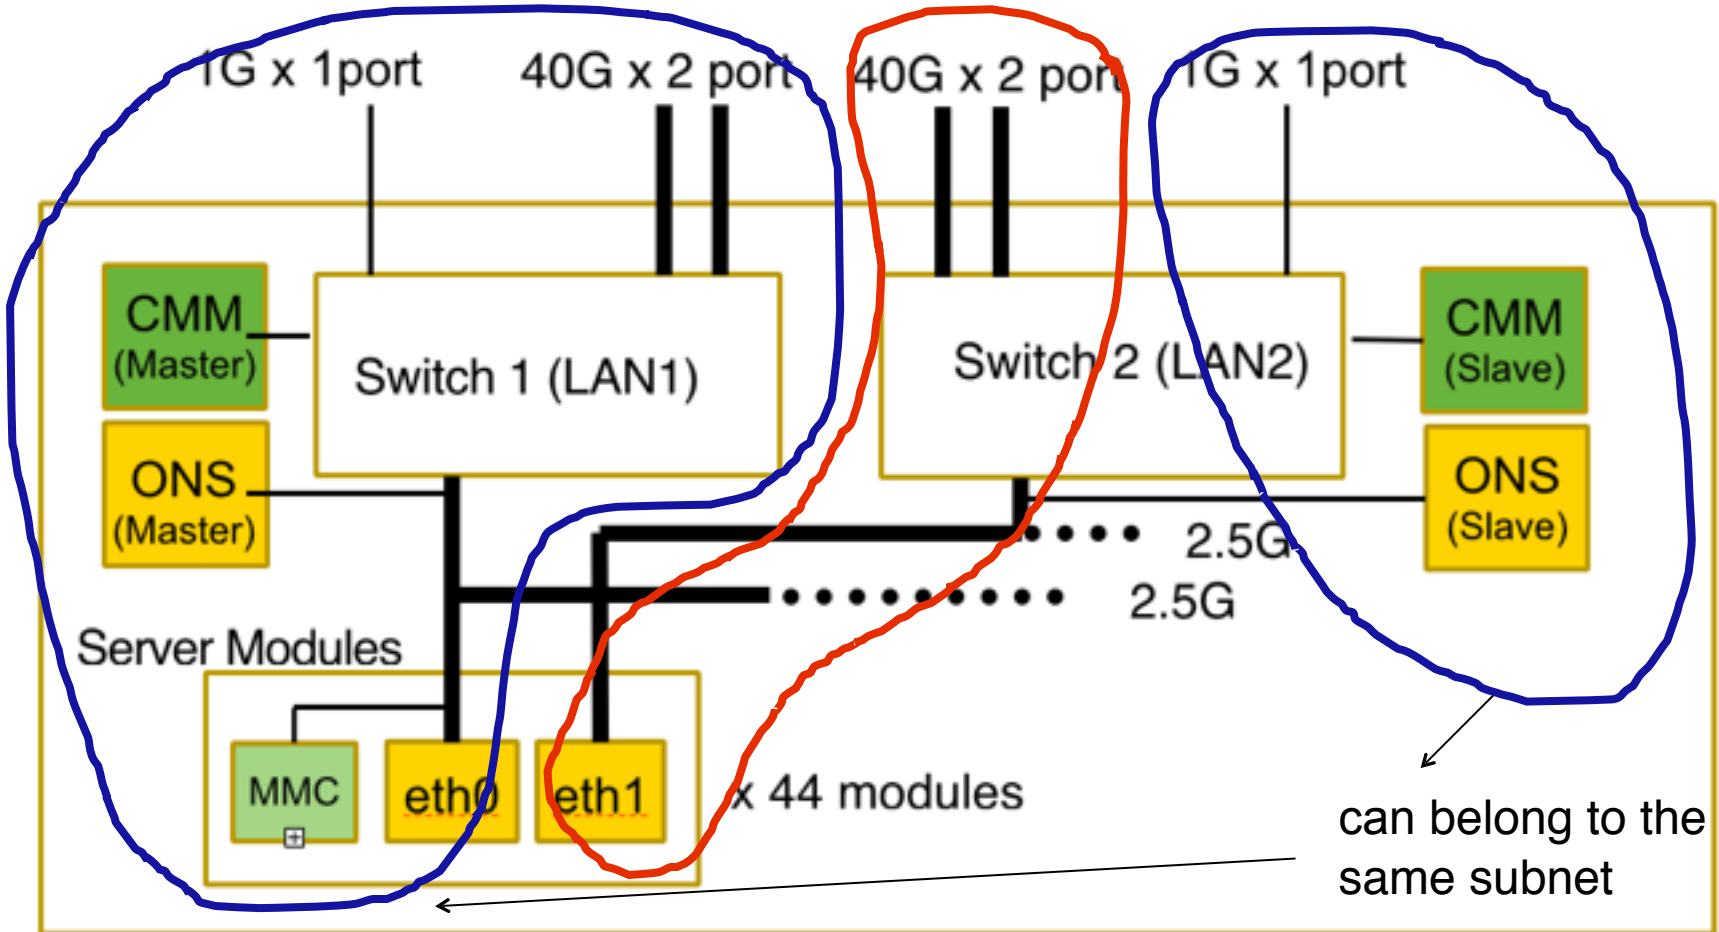
Network Connection in a Chassis

- MMC: (BMC : Base Management Controller) which controls power, boot, etc of each server, only accessible through LAN1.
- CMM: a BMC which controls the chassis function. } Master thru LAN1
- ONS: a control port for each of two switch boards in a chassis. } Slave thru LAN2



VLAN Reconfiguration

- Different VLAN is assigned by default: 1G - VLAN#4092 40G - VLAN#1
- On all chassis, changes 1G-VLAN from #4092 to #1 so that eth0 (OS kernel port) traffic is forwarded Sw1 1G port. On leaf1 and 2, LAN1 1G traffic is forwarded to 40G.
- Since LAN2 is isolated form host network (L3 network), loop problem is resolved.



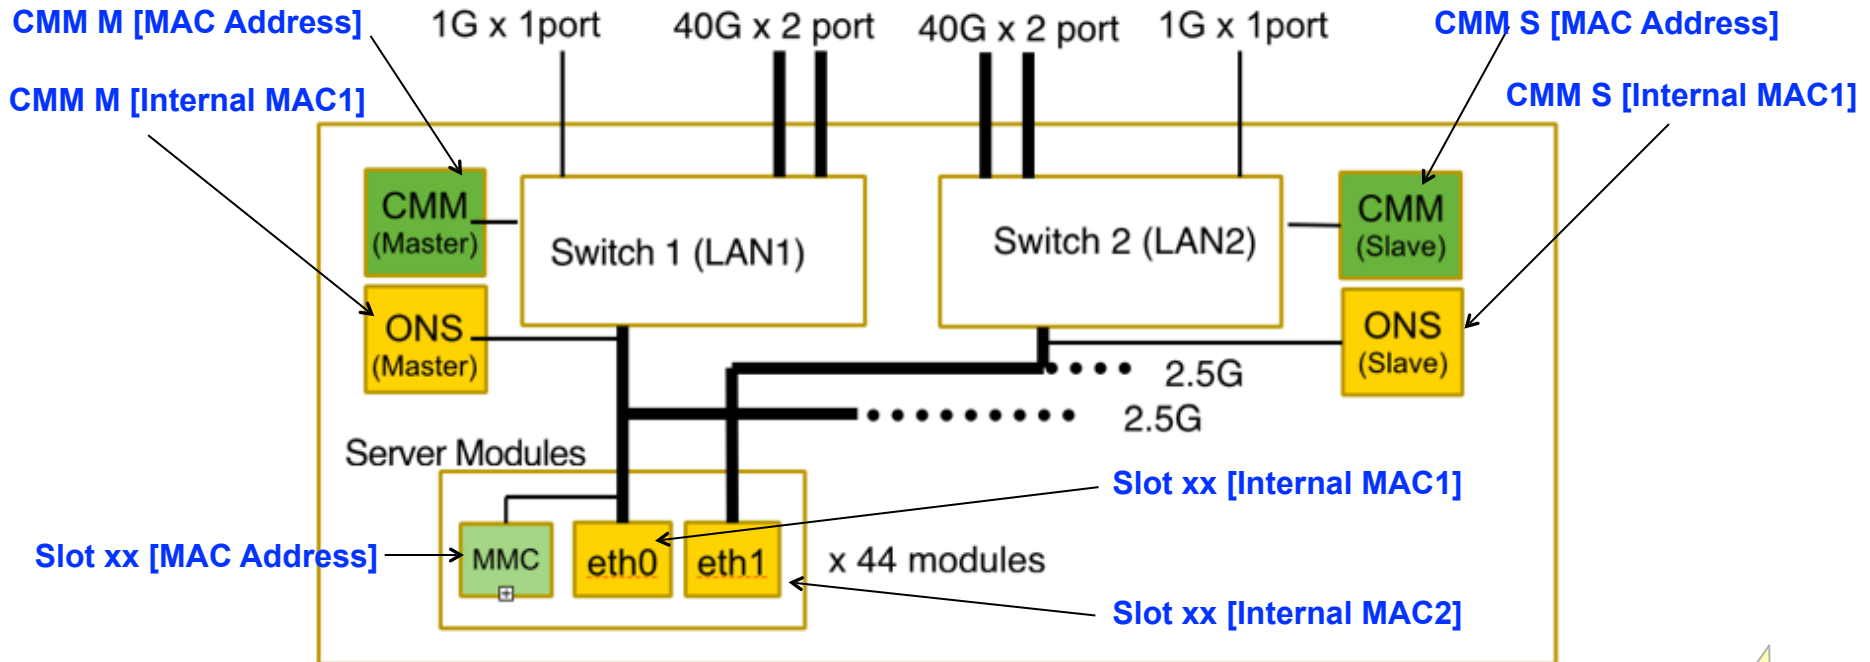
Entities MAC address

- How to read MAC address table acquired through CMM-M

i. MAC address table)

SlotNo	IPAddress	[MAC Address]	F/W	BIOS	GUID	[InternalMAC1]	[InternalMAC2]
CMM M	192.168.0.238	[74:d4:35:1e:83:a0]	1.10	----	80eee18a-ab93-e311-0180-74d4351e83a1	74:d4:35:1e:83:a1	
CMM S	192.168.0.229	[74:d4:35:1e:84:44]	1.10	----	80951243-b293-e311-0180-74d4351e8445	74:d4:35:1e:84:45	
Slot01	192.168.0.220	[74:d4:35:83:77:d4]	1.10	0004	809e90e0-df7d-e311-0180-74d4358377d2	74:d4:35:83:77:d2	74:d4:35:83:77:d3
Slot02	192.168.0.223	[74:d4:35:83:77:f8]	1.10	0004	8072e472-df7d-e311-0180-74d4358377f6	74:d4:35:83:77:f6	74:d4:35:83:77:f7
Slot03	192.168.0.225	[74:d4:35:83:77:d7]	1.10	0004	8006744c-e27d-e311-0180-74d4358377d5	74:d4:35:83:77:d5	74:d4:35:83:77:d6
Slot04	192.168.0.235	[74:d4:35:83:78:58]	1.10	0004	80d06ba6-d07d-e311-0180-74d435837856	74:d4:35:83:78:56	74:d4:35:83:78:57

ii. Relation to the entities in a chassis)




Reconfigure VLAN

1. LAN1 reconfiguration on each server
 - Connect MMC-M of each server with IPMItool through 1G port and issue commands to:
 1. Remove control VLAN #4092 and merge it to VLAN #1
2. LAN1 switch reconfiguration:
 - Connect terminal to microUSB port on the front panel and issue commands to:
 1. Initialize ONS
 2. Merge VLAN #4092 to VLAN#1

Note) Item 1 should be done before item 2, otherwise we lose path to access MMC-Ms.

ONC Access Ports



Win/Mac driver for the USB ports:

<http://www.silabs.com/products/mcu/pages/usbtouartbridgevcpcdrivers.aspx>

IP Addresses mapped by host DHCP

Need to assign 276 IP addresses with DHCP on the host.
At least 9bit class-C subnet is required.

1. IP address assigned to a chassis

CMM(Master)	1
CMM(Slave)	1
ONS(Master)	1
ONS(Slave)	1
MMC	44
eth0 port	44
Total	92

Appendix

