


Enabling FPGAs in Hyperscale Data Centers

J. Weerasinghe¹, F. Abel¹, C. Hagleitner¹, A. Herkersdorf²

¹IBM Research – Zurich Laboratory
²Technical University of Munich

- Background
- Motivation
- Hyperscale Data Center
- Network-attached FPGAs
- OpenStack Integration
- Conclusion



Heterogeneous workloads are moving into the cloud.
How do we accelerate them?



Cloud(Public/Private)



Accelerate Workload Processing in the Cloud

- Research
- Educational
- Testing

- FPGA-accelerated APIs

- Improved user experience (e.g. BING)




FPGAs as a Service

Acceleration as a Service

IaaS

PaaS

SaaS



Cloud(Public/Private)



Scalability



Optimal
Utilization

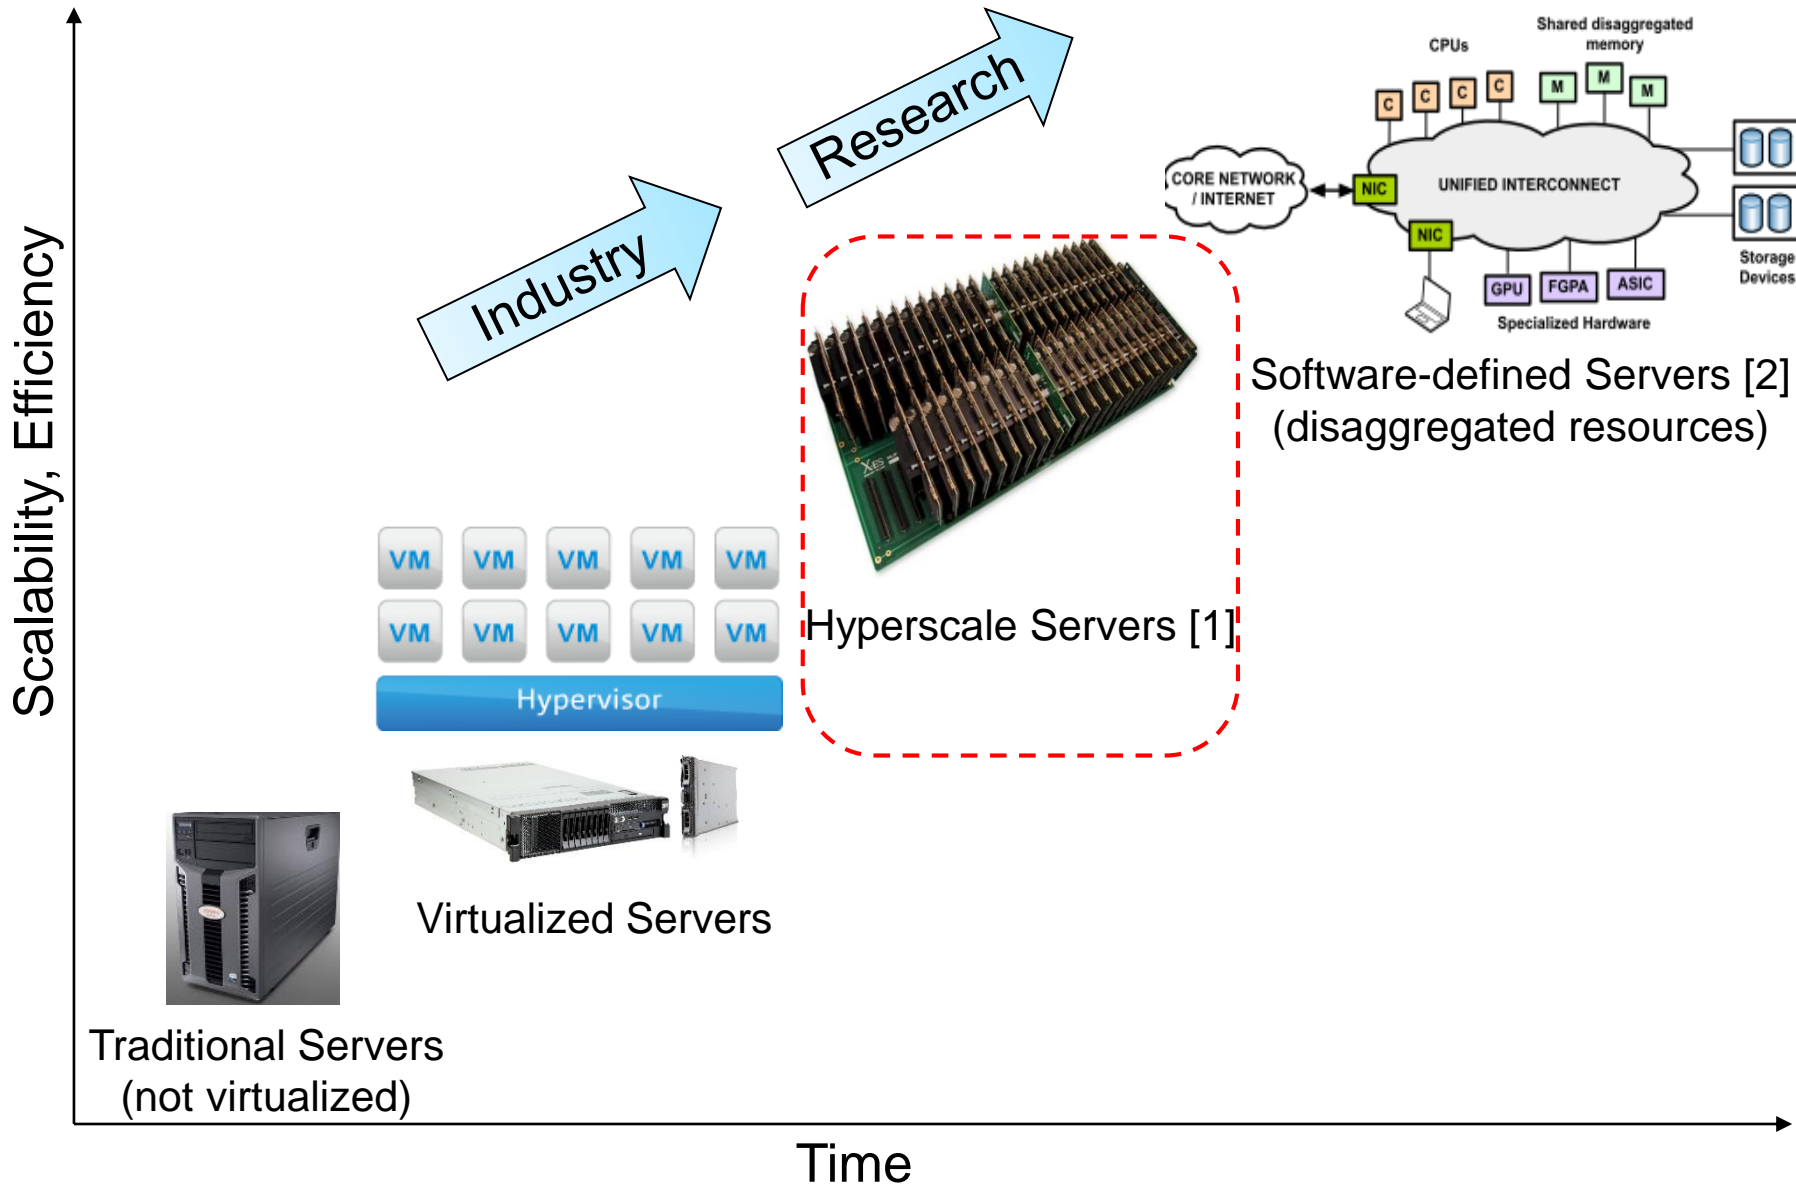


Fail-in-
Place



Workload
Migration


Server/Cloud Data Center Trends



[1] Cray and X-ES Microserver

[2] Network support for resource disaggregation in next-generation datacenters (UC Berkeley & Huawei, 2013) © 2015 IBM Corporation

Hyperscale Data Center (e.g. IBM DOME uServer)

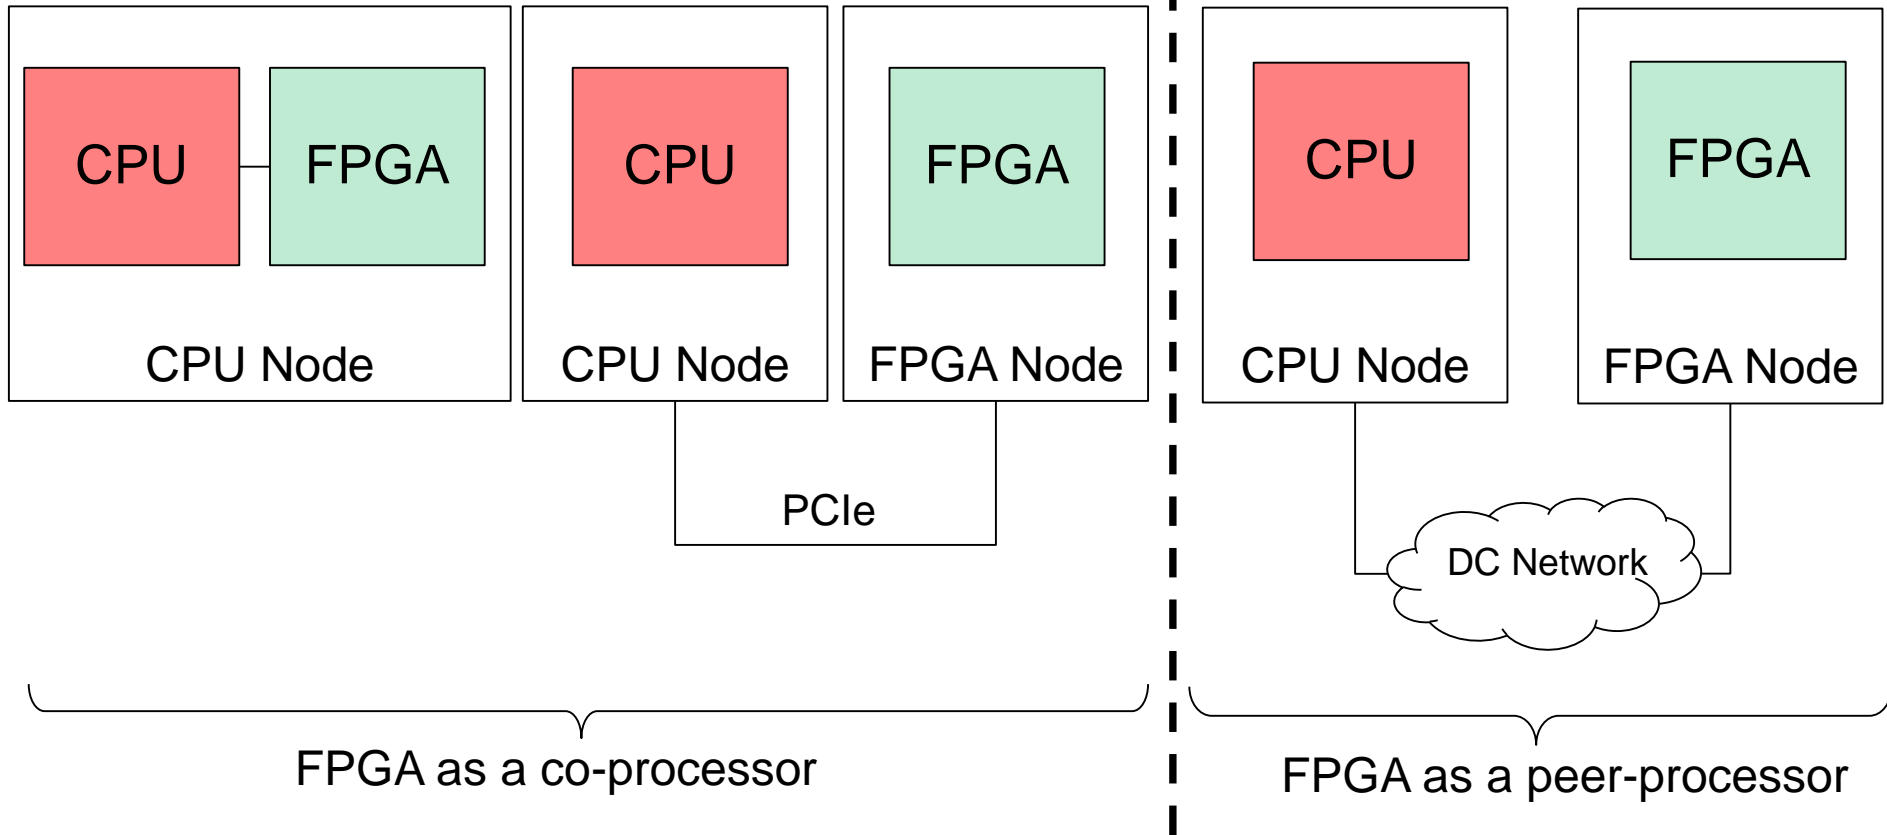


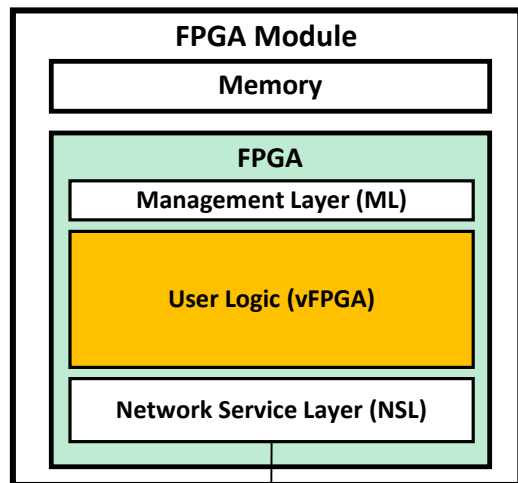
Compute module with
up to 48GB DRAM

128 Compute modules
in a 2U chassis

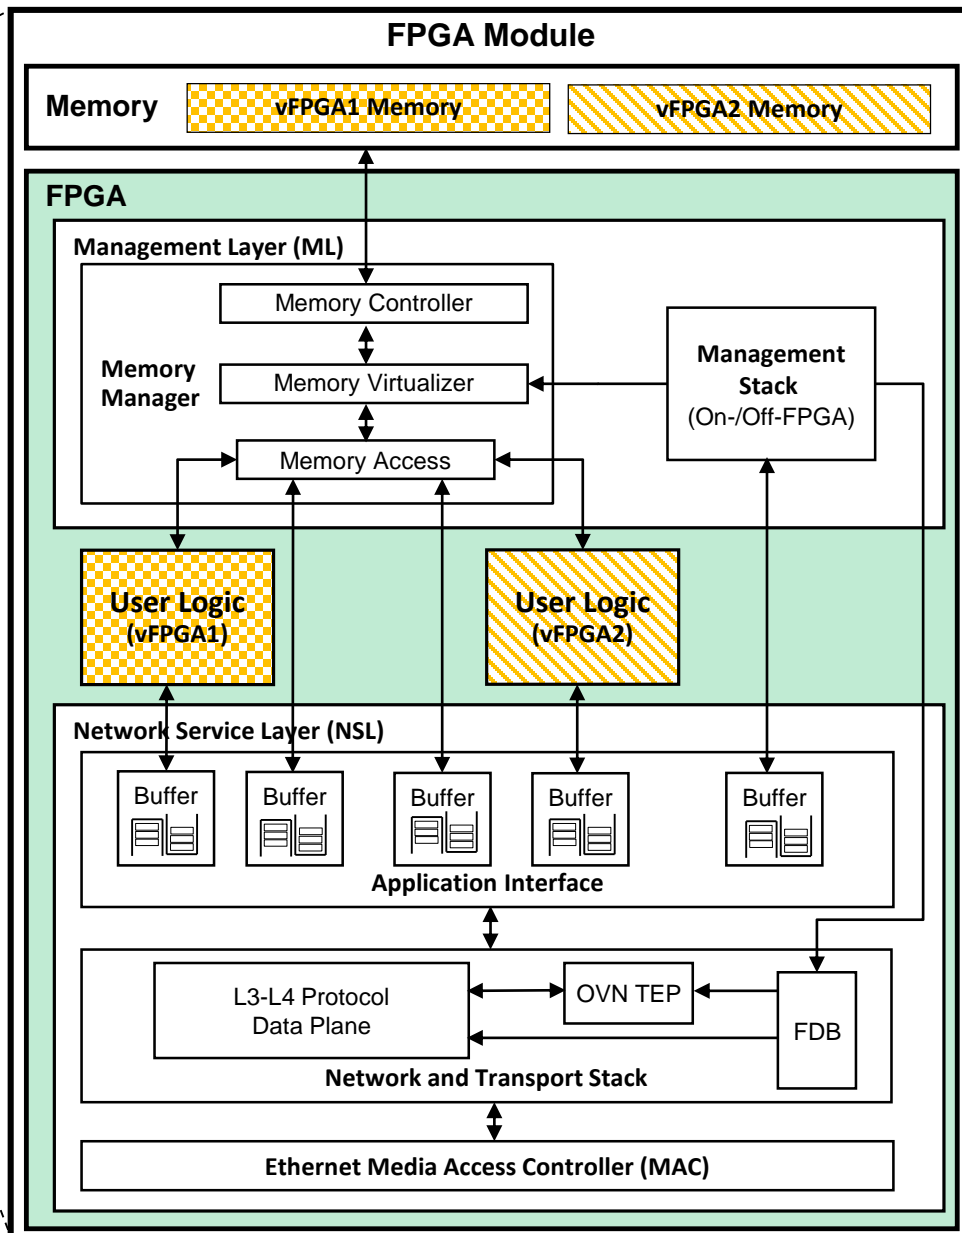
8K Compute modules &
400TB DRAM in 4 racks

FPGA-CPU Attachment Options

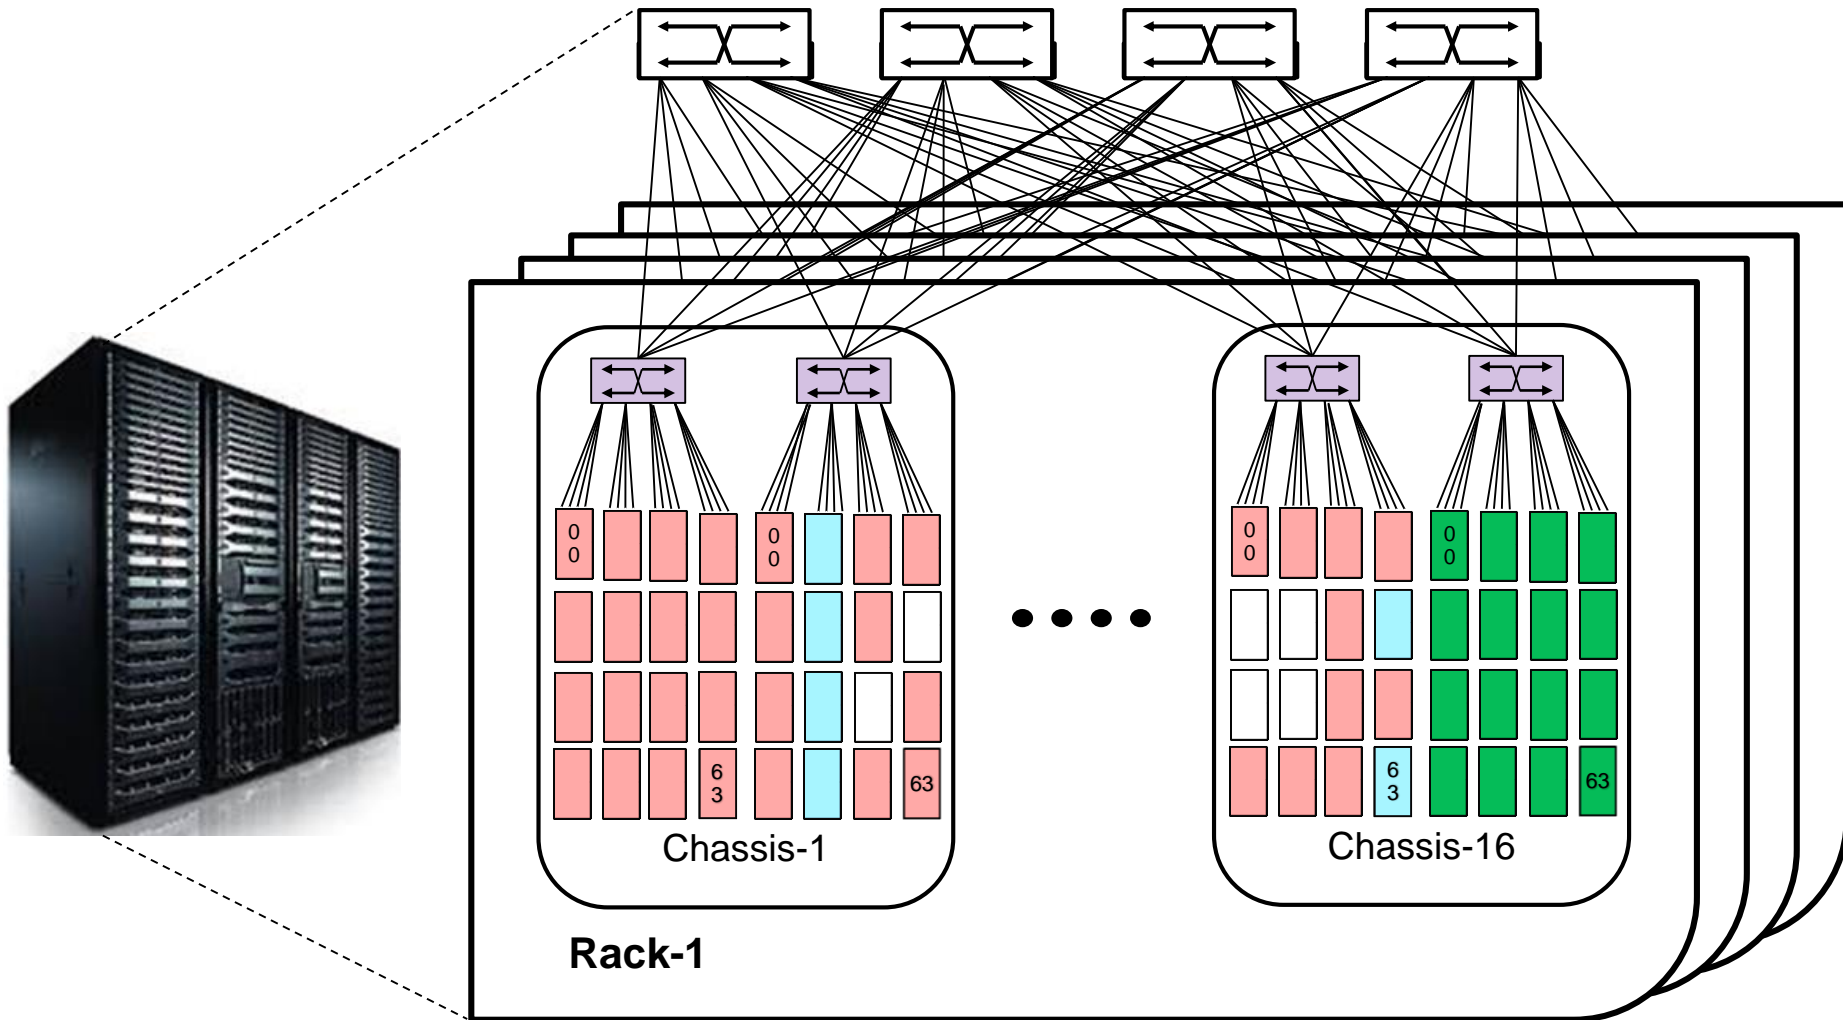




PRR: Partially Reconfigurable Region, FDB: Forwarding DB, OVN: Overlay Virtual Network, TEP: Tunnel Endpoint



Hyperscale Data Center w/ Network-attached FPGAs



Compute Module Storage Module FPGA Module Other Modules

Networking Module

Top-of-Rack Switch

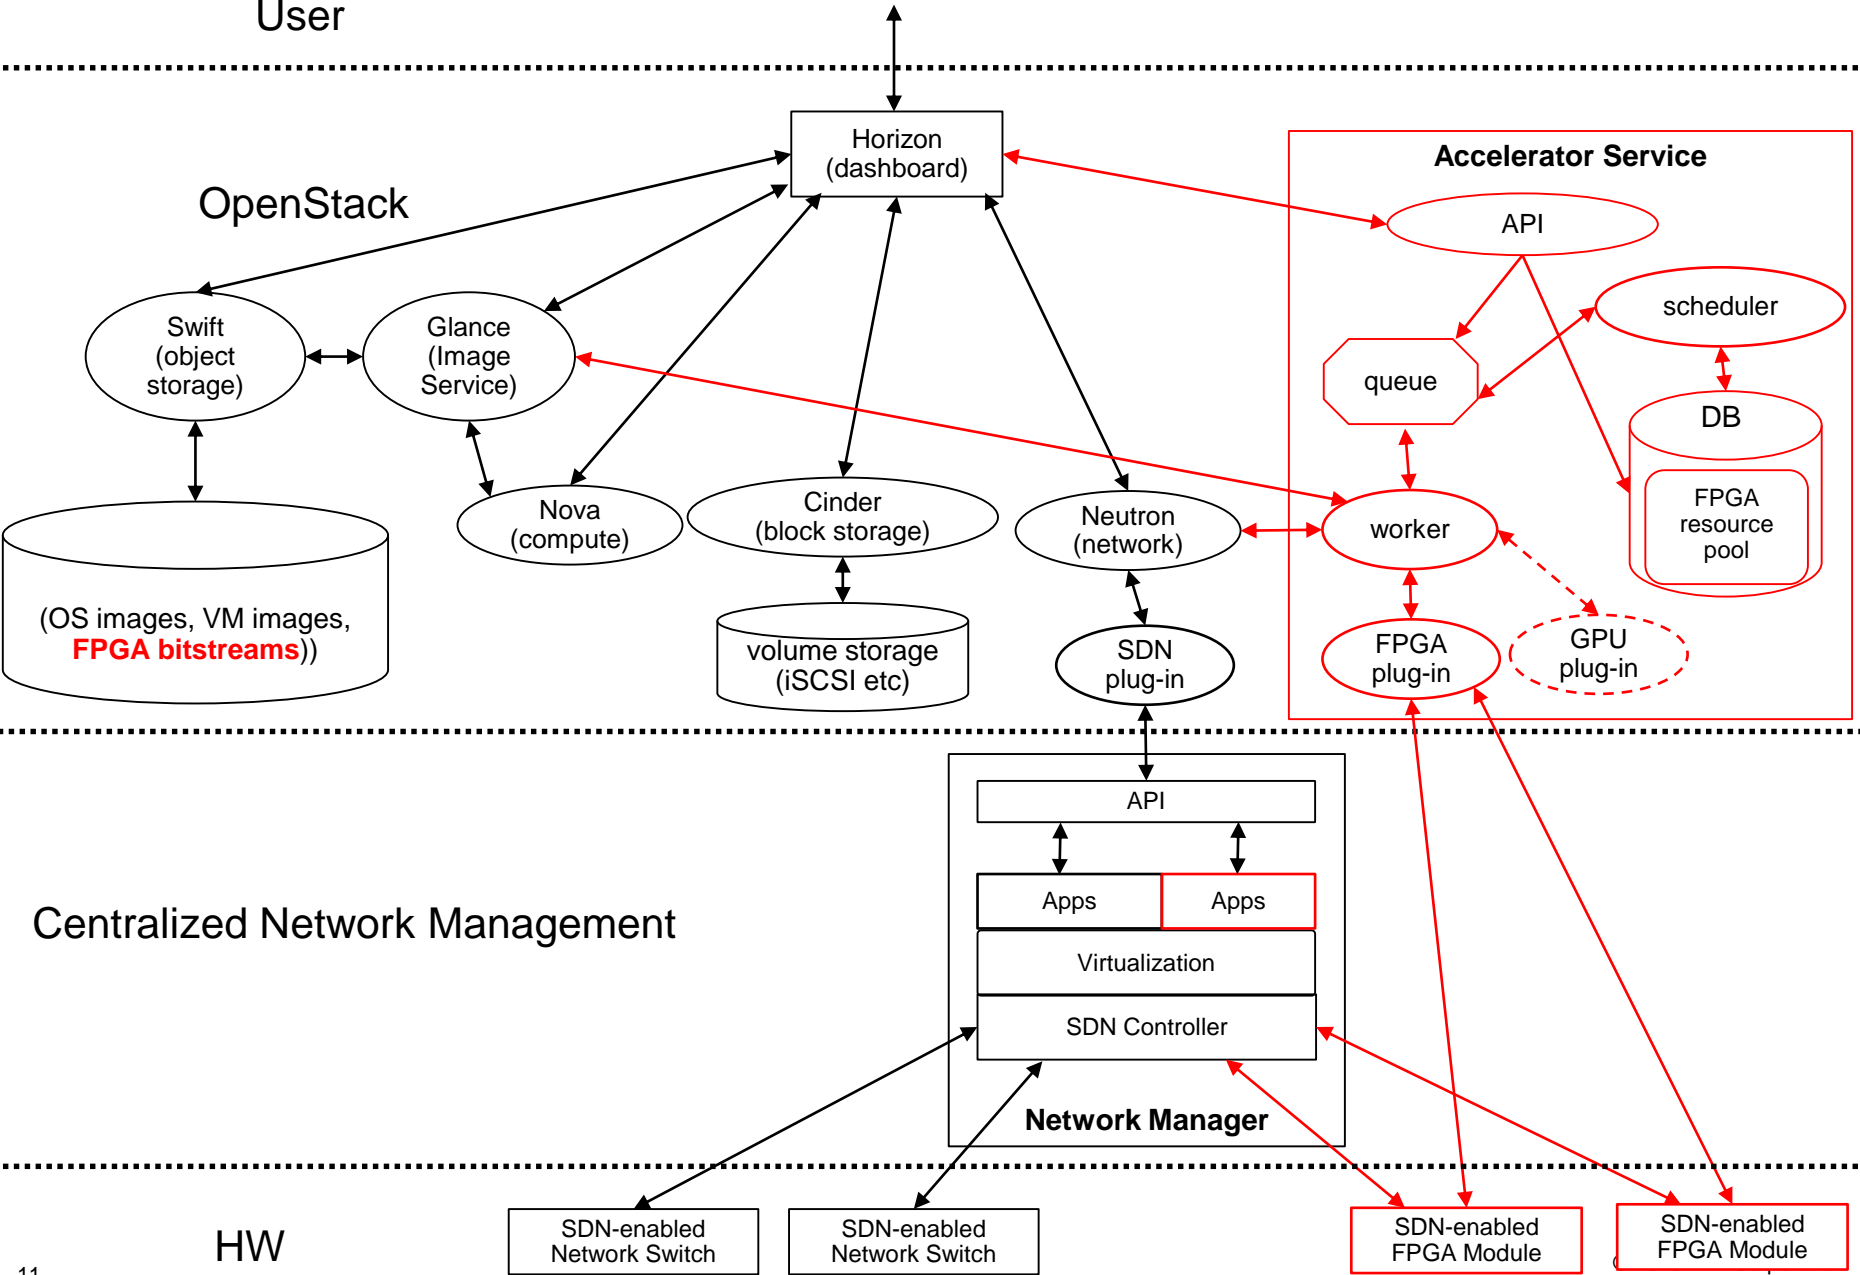
User

OpenStack

Accelerator Service

Centralized Network Management




HW




vFPGA: Rent Your FPGA in the Cloud



IP Address: 10.10.1.8
Memory Capacity: 8GB
CPU Cores: 4



IP Address: 10.10.1.9
Memory Capacity: 8GB
Resources:
CLB: 100000
DSP: 800
BRAM: 400



Streaming?

Compute Module


SW Application

FPGA Fabric

Compute Module

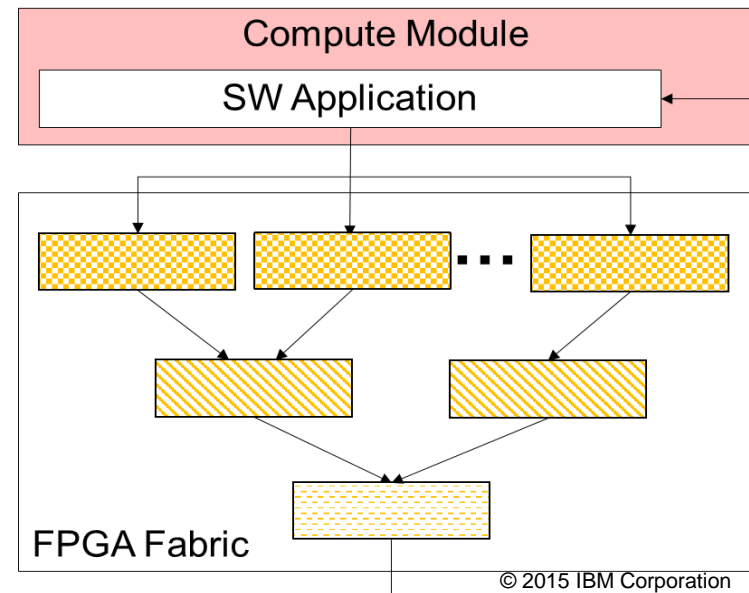
SW Application

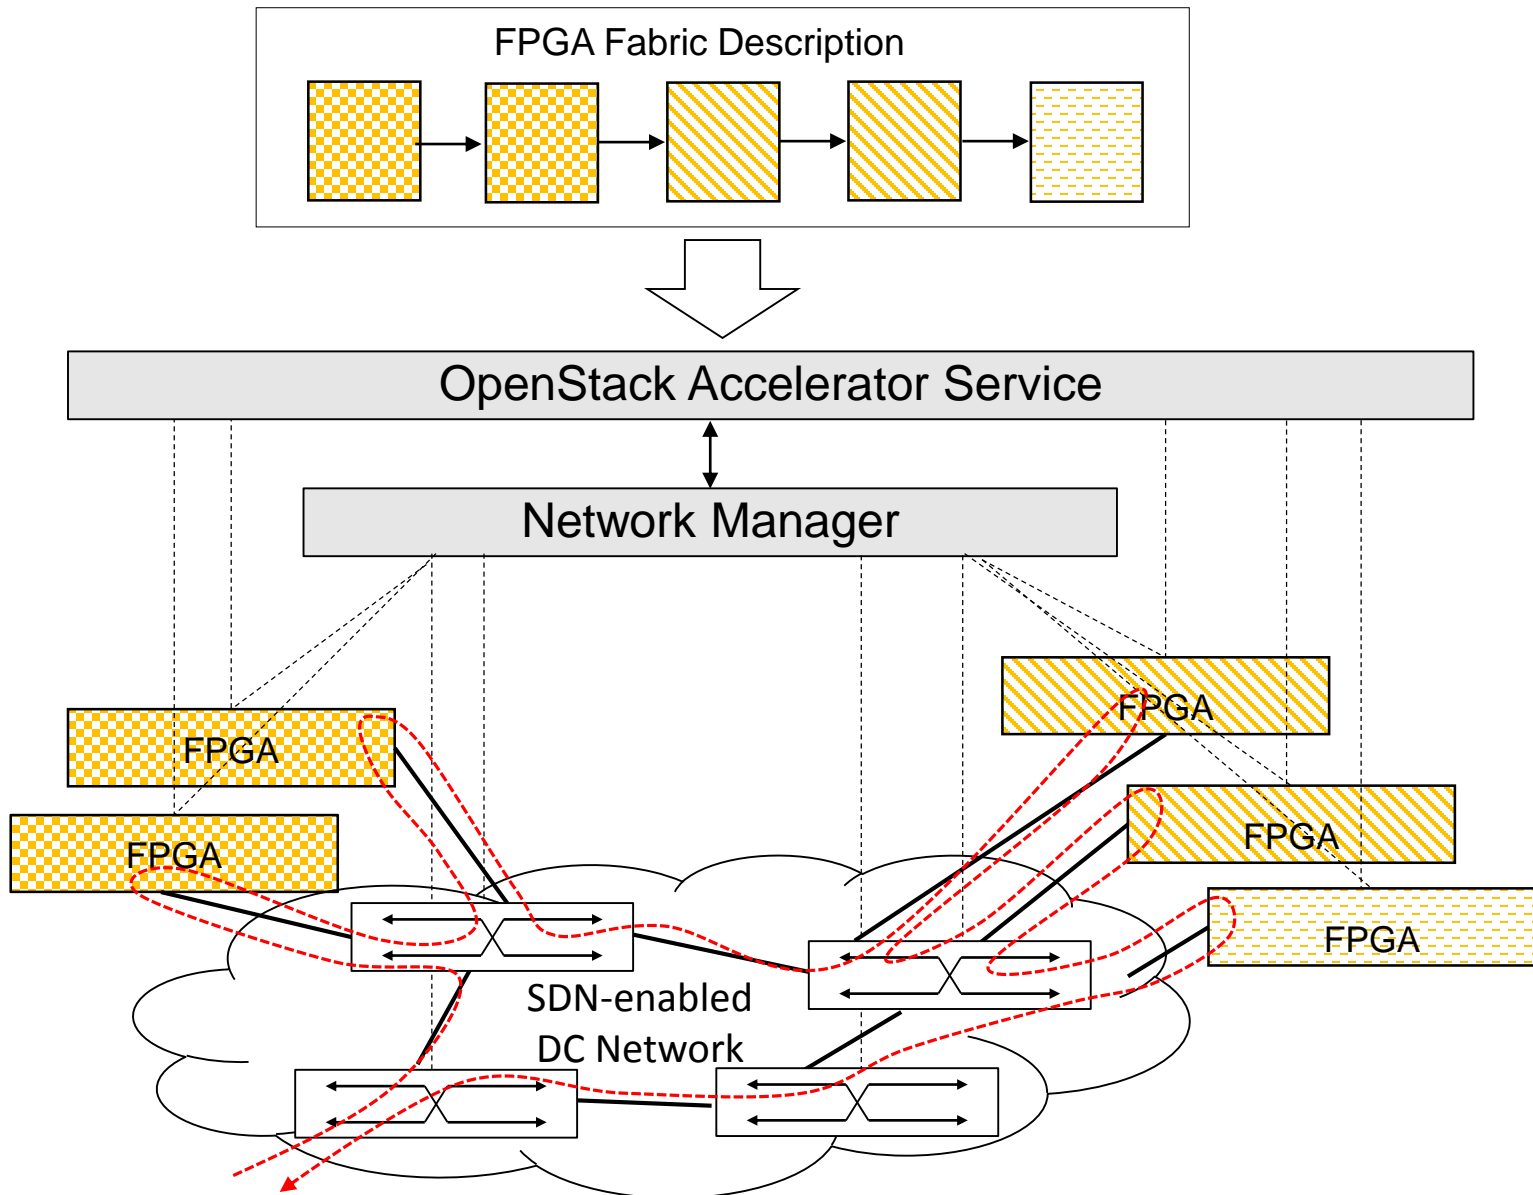
FPGA Fabric

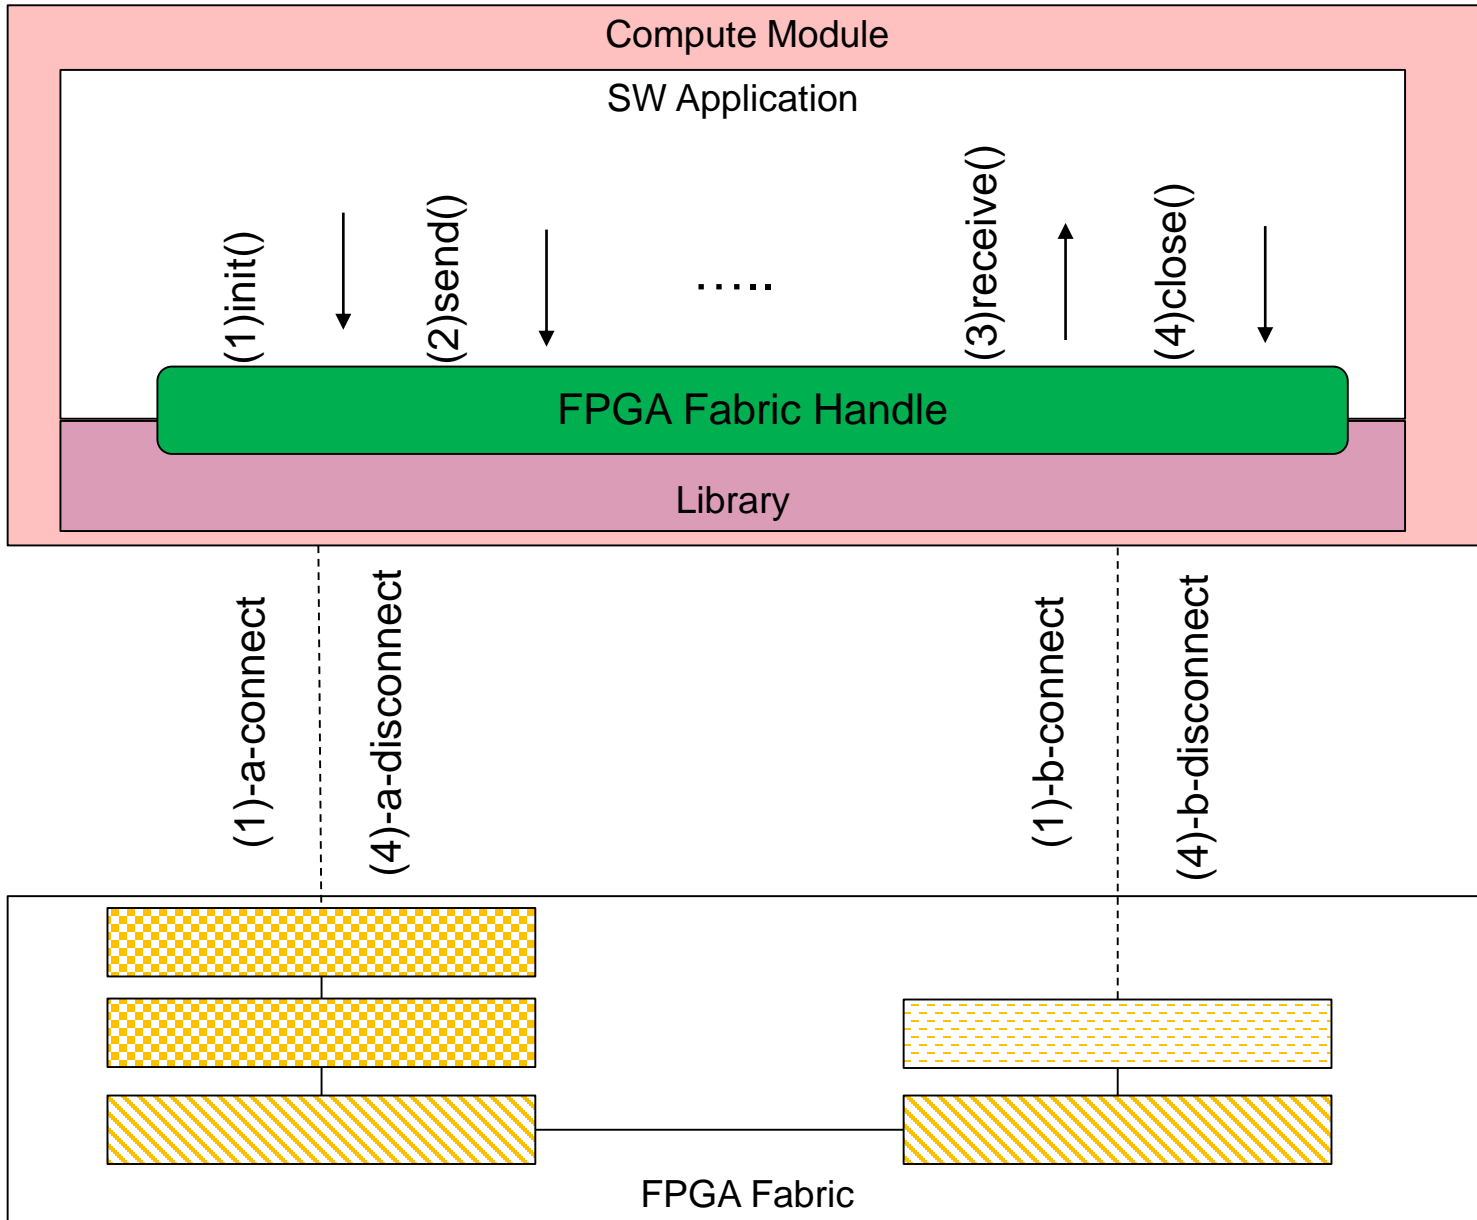


FPGA Resource Pool

MapReduce?
Neural Networks?







Per Rack Compute Performance

Per Rack	Hyperscale Server [1]	Hyperscale FPGA [2]
Modules	2048	2048
Cores	24586	4096
CLBs(10 ⁶)	-	71
Memory(TB)	100	100
TFLOPS(32 bit)	442	958

[1] FreeScale T4240 SoC-based

[2] Xilinx Zynq 7100 SoC-based

- FPGAs must be deployed in data centers
- FPGAs must be made available to the cloud users
- Large-scale deployment requirements
 - A scalable infrastructure
 - A change in the CPU-FPGA interface
 - A management framework
- We propose
 - Network-attached FPGAs in hyperscale data centers
 - OpenStack accelerator service to provision FPGAs in the cloud

***Make the FPGA Free from CPU
and
Go Large Scale***

THANKS

