



# Cloud FPGA



[bit.ly/cloudfpga](https://bit.ly/cloudfpga)



## Lecture: Recent Cloud FPGA Advances – IBM

Prof. Jakub Szefer

Dept. of Electrical Engineering, Yale University

EENG 428 / ENAS 968


Cloud FPGA




Share:  
[bit.ly/cloudfpga](http://bit.ly/cloudfpga)

EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019


# FPGAs in Data Centers



- Many companies have been exploring with FPGAs in data centers
  - Use of FPGAs has gone from accelerators to publicly available Cloud FPGAs
  - Major cloud providers: Amazon, Microsoft, etc.
- Other deployments of FPGAs in cloud data centers include **cloudFPGA project from IBM**
- IBM operates FPGAs in a public Cloud of the Zurich Research Laboratory (ZRL) referred to as Zurich Yellow Cloud 2 (ZYC2)
  - Ongoing research project
  - Similar to TACC and Catapult
  - Interesting features such as direct access of FPGAs to data center network




# IBM Netezza




Share:  
[bit.ly/cloudfpga](https://bit.ly/cloudfpga)





- Netezza offered a **data warehouse appliance (DWA)**
- A pre-installed and pre-optimized hardware and software for big data analysis
  - Not a publicly available cloud service
  - Instead a system sold to companies to use internally
- Part of the hardware were the S-Blades
  - An independent server that contains multi-core Intel-based CPUs and Netezza's proprietary high-throughput FPGAs
- Discontinued around 2019
  - Could not easily adapt to cloud computing paradigm



IBM Netezza 100 system, a single-chassis, single-host, system with one S-Blade

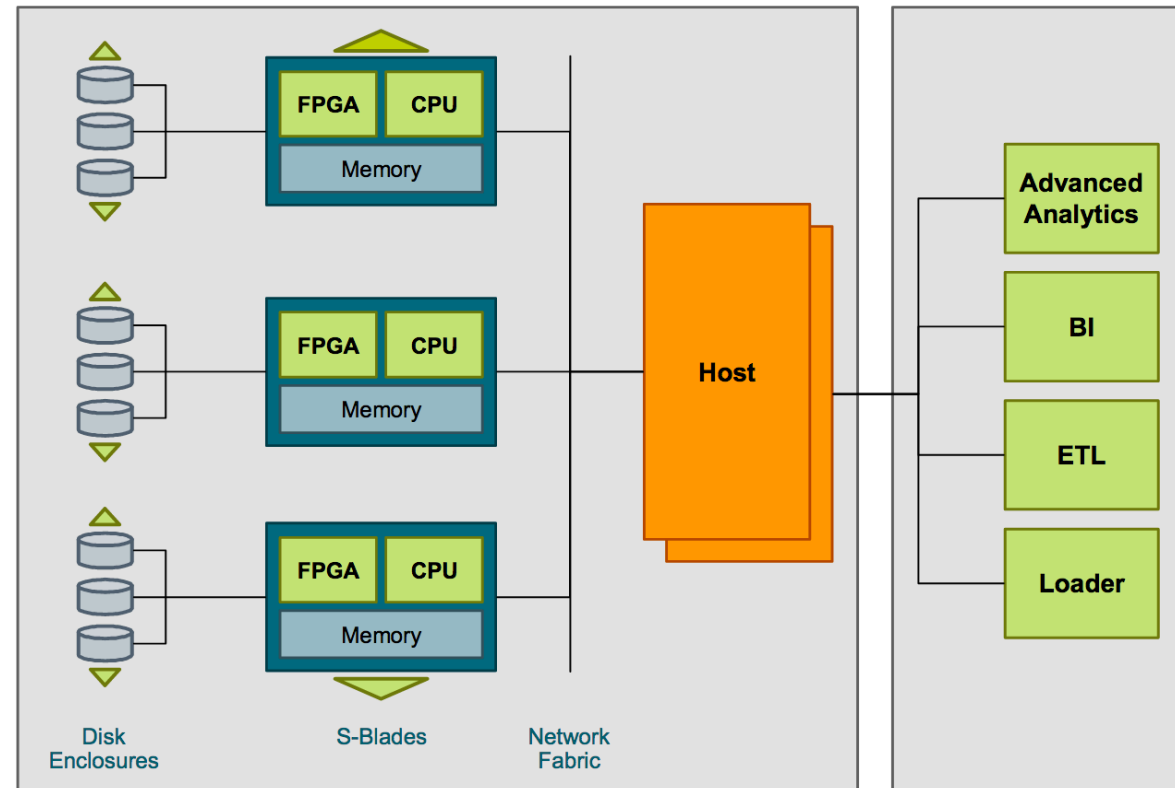


IBM PureData System for Analytics N1001 systems, with 4 to 14 S-Blades


# IBM Netezza FGAs



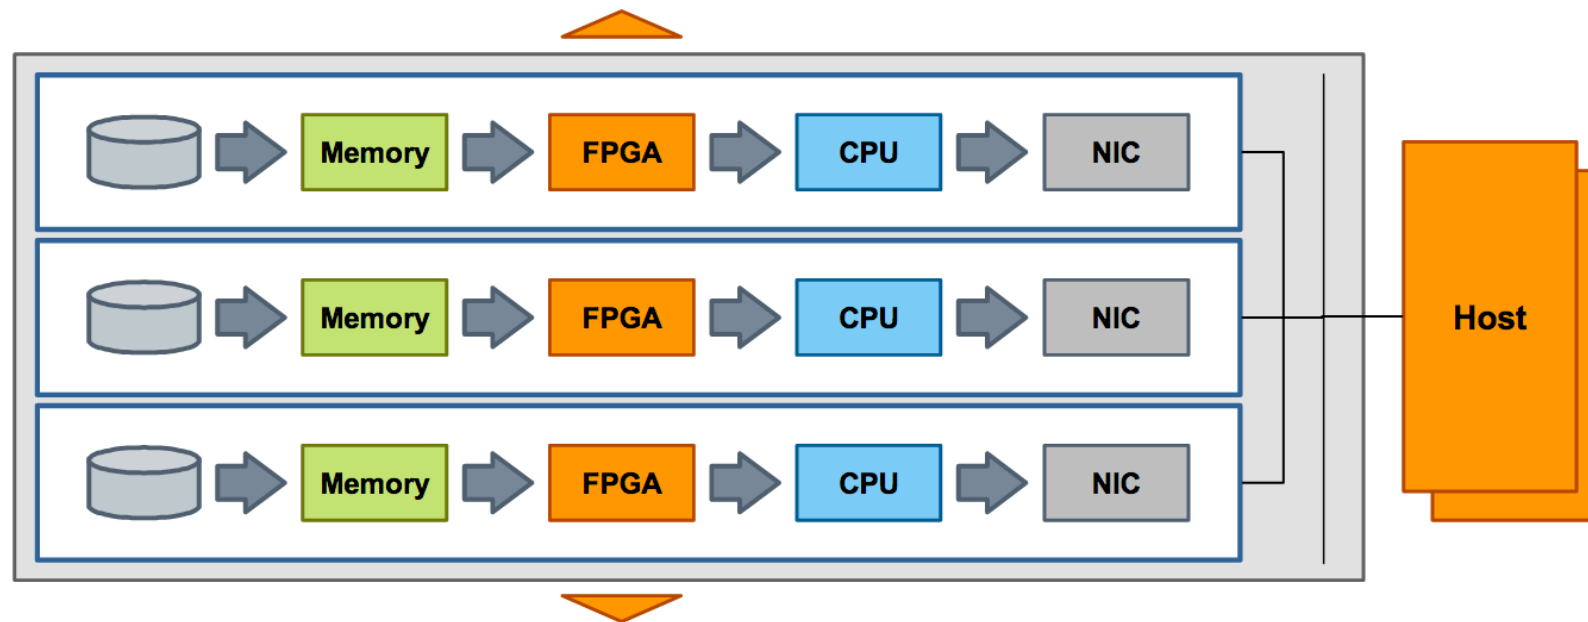
- Netezza used a special field programmable gate array (FPGA) co-processor to speed up a heavily modified PostgreSQL database
- There were eight FGAs on the accelerator blade - one for each x64 core on the paired server blade
  - FGAs were used to speed up the filtering of data moving off storage before being passed on to the database software
  - And doing complex sorting and joins of database tables as part of analytical routines




# IBM Netezza FPGA S-Blades



- Purpose of FPGA is to pre-process data coming in from memory
  - Filter or do basic functions as the speed of memory
  - Remainder of operations is done by CPU
- Effectively a custom, hardware data base processing engine



# “cloudFPGA” from IBM



Share:  
[bit.ly/cloudfpga](https://bit.ly/cloudfpga)




*The **cloudFPGA** project aims to provide FPGAs in cloud computing data centers in a disaggregated manner*

- A platform to deploy FPGA at large scale in data centers
  - Integrates FPGAs at the drawer/chassis layer
  - Combines passive and active water-cooling
  - Provides high density, energy efficiency and reduced costs
  - Fits 1024 FPGAs per data center rack
- Disaggregates FPGAs from the servers
  - FPGAs connect to the data center network over 10/40 Gb/s Ethernet
  - FPGA cards are stand-alone resources, independent of the number of servers
  - Makes FPGAs plentiful in data centers, users can rent and link them in any type of topology



# Enabling FPGA as a Service



Accelerate workload processing in the cloud

- Research
- Education
- Evaluation

- Accelerated APIs & APPs (e.g. Genomics)

- Improved user experience (e.g. BING)


FPGA as a Service

Acceleration as a Service

IaaS

PaaS

SaaS




Share:  
[bit.ly/cloudfpga](http://bit.ly/cloudfpga)

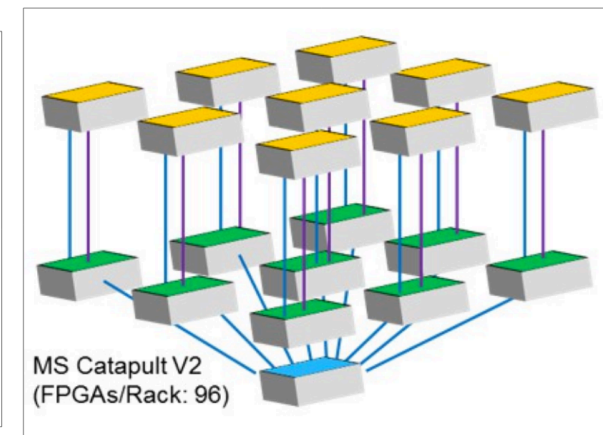
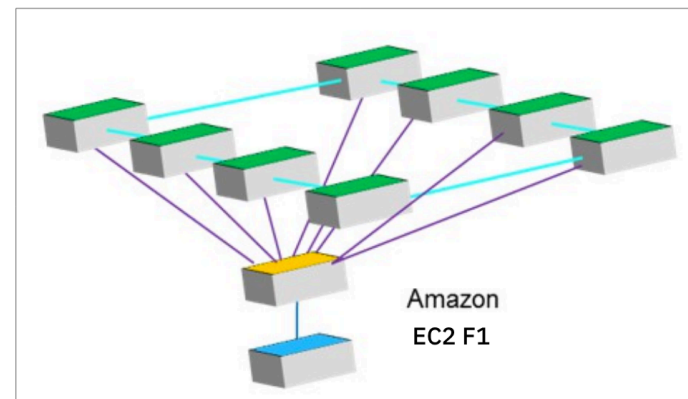
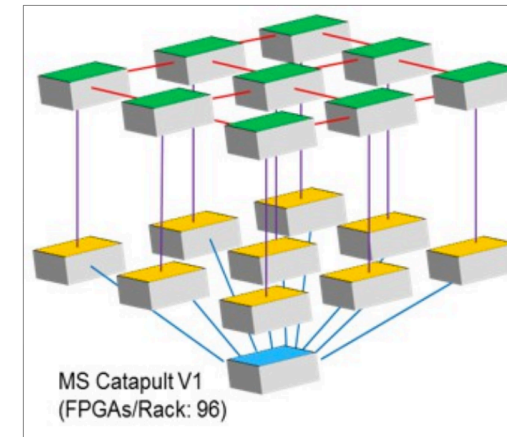

EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019

Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>

# Comparison of FPGA Deployments in the Cloud



- Different companies and research projects take different approach to developing Cloud FPGA architectures
  - Attach FPGA to CPU
  - FPGA between network and CPU
  - FPGA attached directly to the network




Share:  
[bit.ly/cloudfpga](https://bit.ly/cloudfpga)


EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019

Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>

# Typical Use of FPGAs as CPU Accelerators



- Most commonly FPGAs are deployed as accelerators used by the CPU
  - Connect to CPU on system bus
  - Connect to CPU via PCIe
- Limitation in terms of FPGAs per CPU
- Accelerator FPGAs depend on CPU for access to any resources
  - Memory, network, devices, etc.




Share:  
[bit.ly/cloudfpga](http://bit.ly/cloudfpga)

EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019

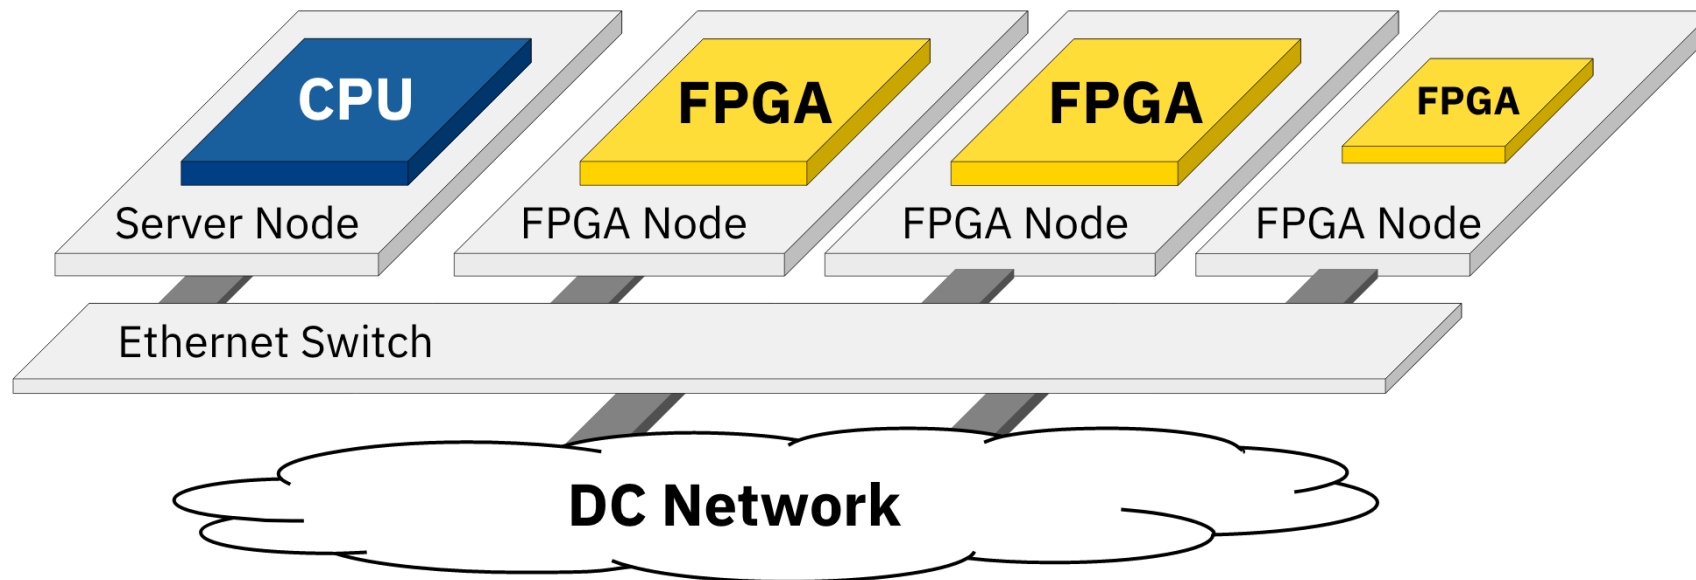
Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>



# Towards Disaggregation of FPGAs



- Disaggregation of FPGAs is the idea to make FPGAs separate from CPUs
  - Each FPGA is a node on network
  - Can be accessed by any CPU or any other FPGA
  - Similar in idea to remote DRAMs, or remote GPUs




Share:  
[bit.ly/cloudfpga](https://bit.ly/cloudfpga)


EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019

Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>


# Towards Disaggregation of FPGAs



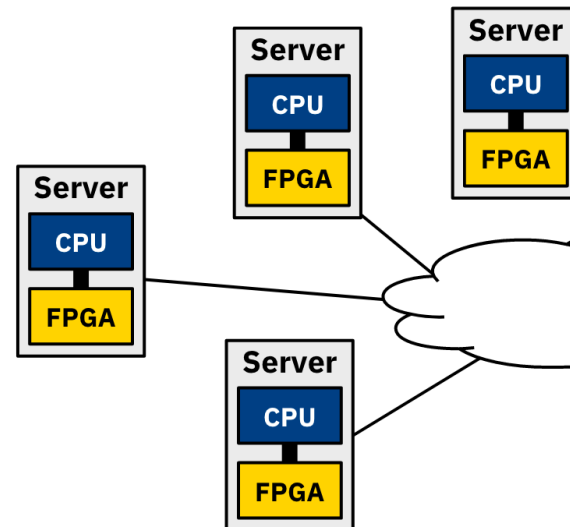
FPGA as a Co-Processor



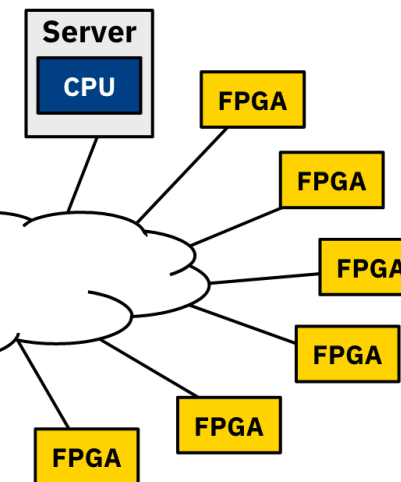
FPGA as a Peer-Processor



CPU-Centric Deployment



FPGA-Centric Deployment




Share:  
[bit.ly/cloudfpga](http://bit.ly/cloudfpga)


EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019

Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>

# Network-Attached FPGAs



- To realize disaggregation of FPGAs, a network-attached FPGA is needed
- FPGA needs to be stand-alone device
  - Use direct network connection
  - Need way to program FPGA

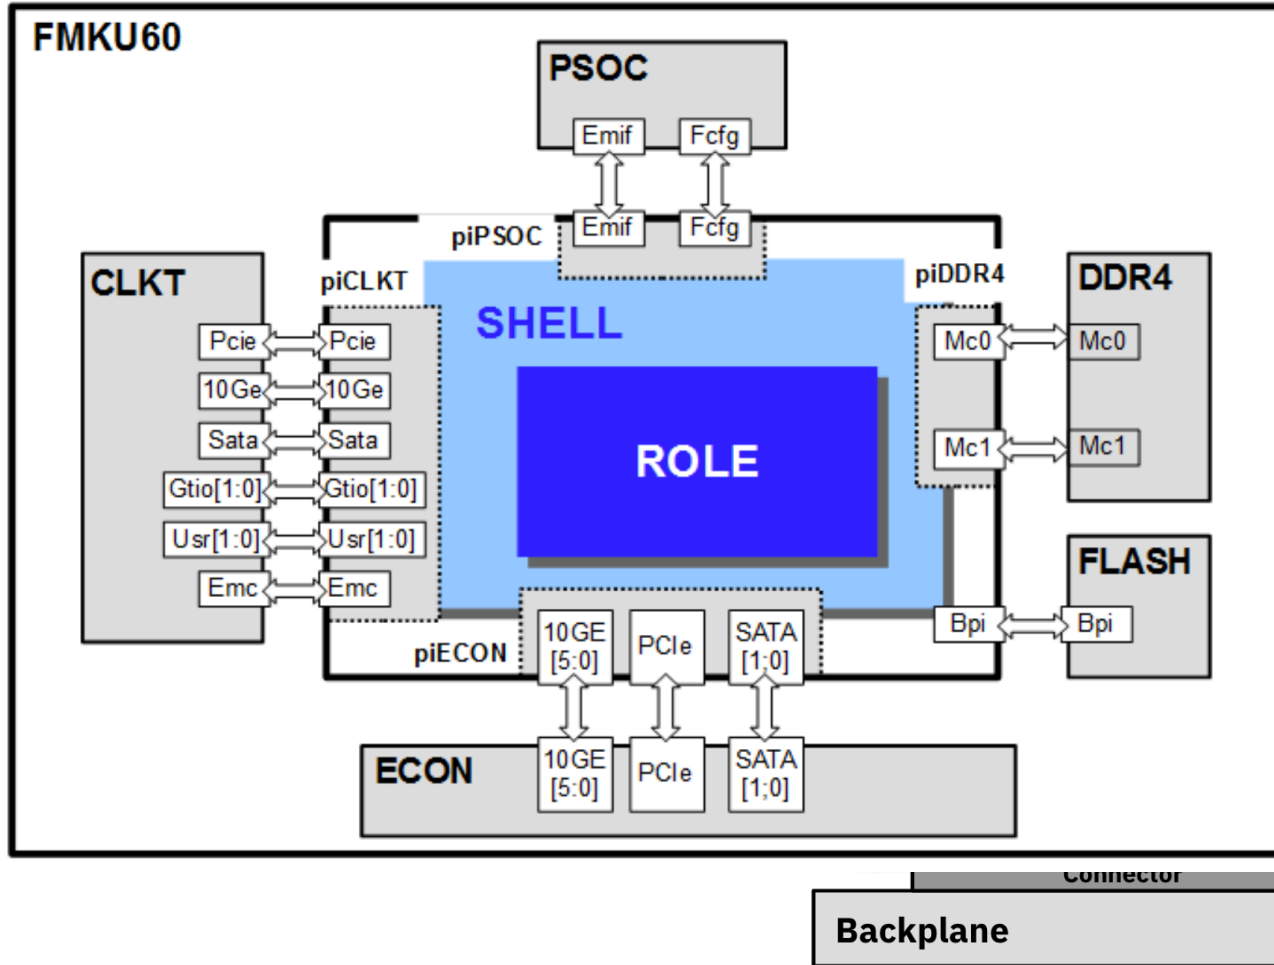


Early Xilinx Kintex UltraScale KU060 used by cloudFPGA project from IBM




Share:  
[bit.ly/cloudfpga](https://bit.ly/cloudfpga)

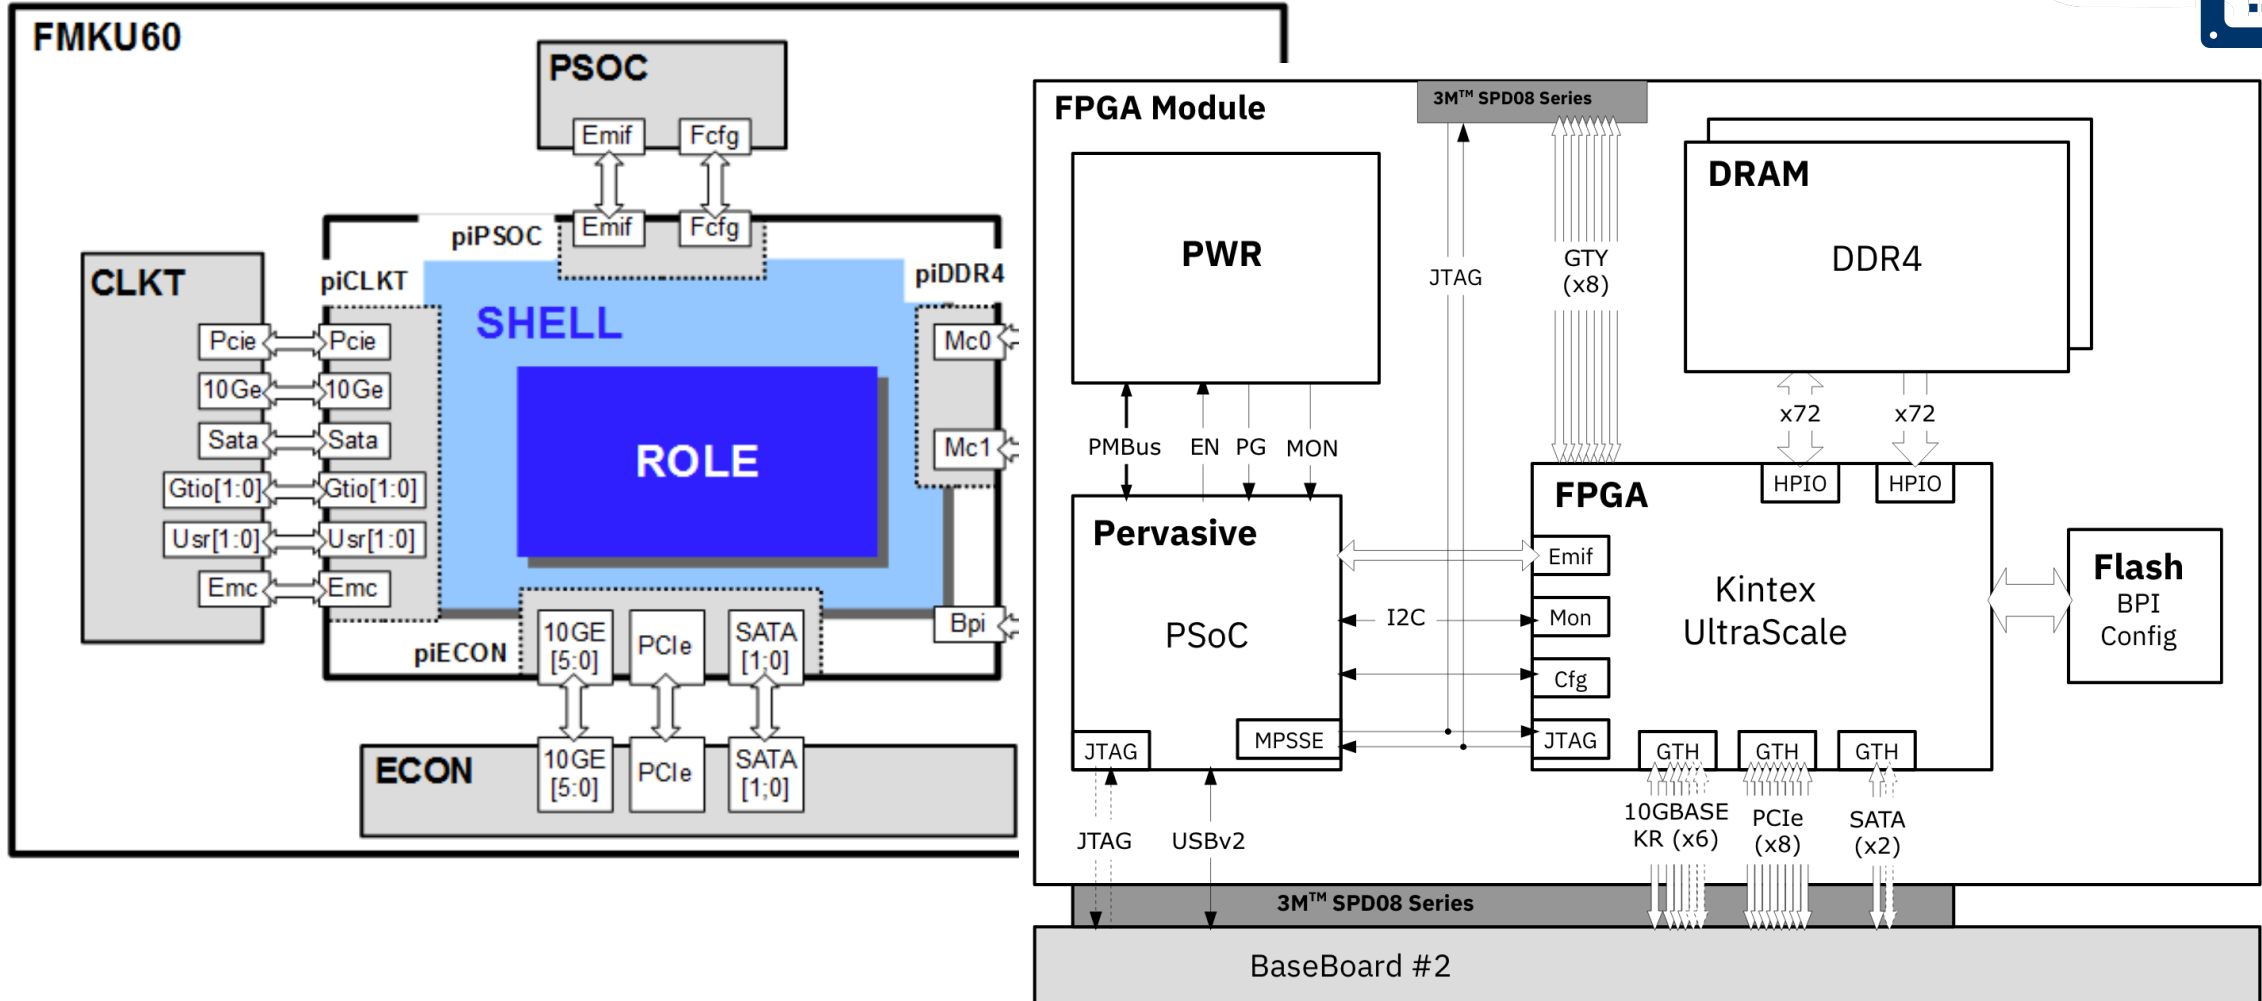

# Shell and Role inside the FPGA



- IBM's Shell and Role are same to Amazon's Shell and CL



# Shell and Role inside the FPGA




Share:  
[bit.ly/cloudfpga](https://bit.ly/cloudfpga)

EENG 428 / ENAS 968 – Cloud FPGA  
 © Jakub Szefer, Fall 2019


Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>



# Network-Attached FPGA Carrier Module



- Through use of custom carrier module, 32 FPGAs are integrated into one module by IBM
  - Including networking switch
  - And USB for programming the FPGAs




Share:  
[bit.ly/cloudfpga](https://bit.ly/cloudfpga)


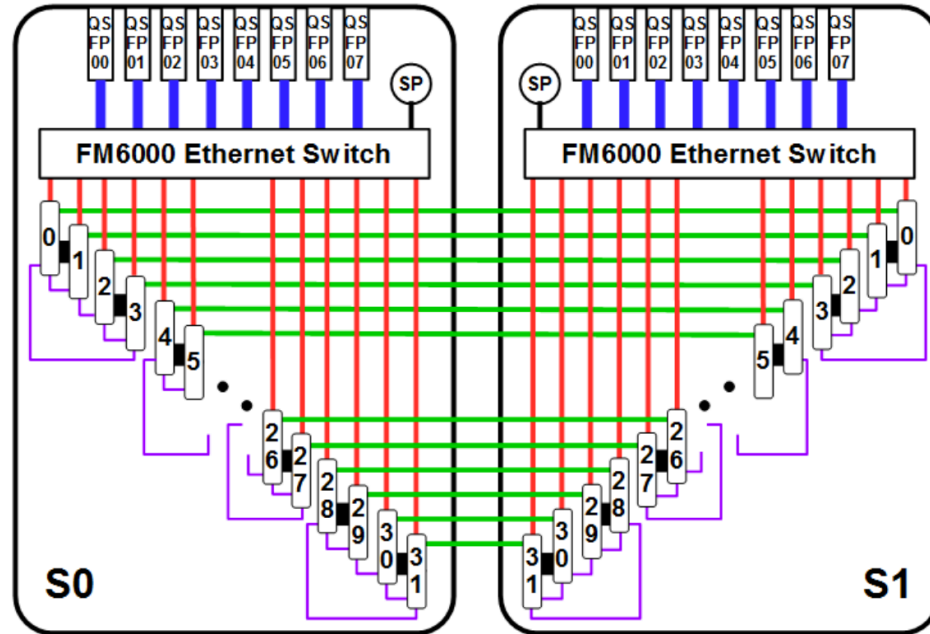
EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019

Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>

# FPGA Carrier Modules in Server Rack Chassis



- Two carrier modules let IBM integrate 64 FPGAs per server rack chassis



Figurative picture

## Legend (per slice):

- [==] x8 40GbE up links (320 Gb/s)
  - [--] x32 10GbE FPGA-to-Switch links (320 Gb/s)
  - [--] x32 10GbE redundant links
  - [--] x32 10GbE FPGA-to-FPGA links
  - [■] x16 PCIe x8 Gen3
  - SP x1 Service Processor
- } Balanced (i.e. no over-subscription) between the north and south links of the Ethernet switch



Share:  
[bit.ly/cloudfpga](https://bit.ly/cloudfpga)

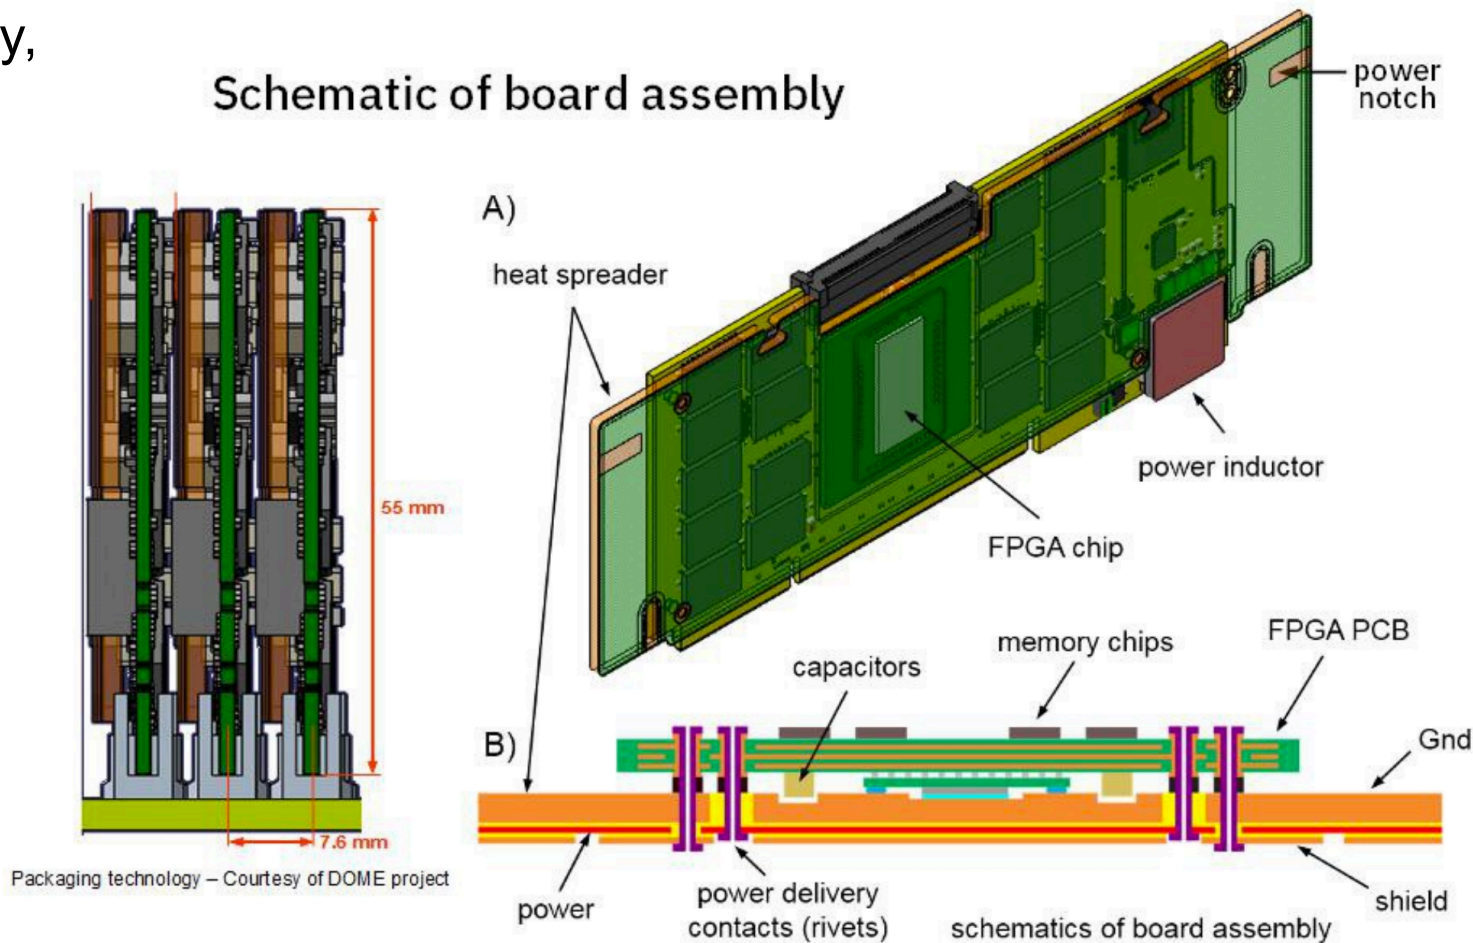
EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019

Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>

# Solving Cooling Issues with High-Density FPGAs




- To achieve high FPGA density, passive cooling and water cooled servers are used
- Each FPGA board does not require a fan nor heat sink
- Use thermally conductive heat spreader...






# Solving Cooling Issues with High-Density FPGAs



- ...heat spreader is connected to a cooling rail
- Use water cooling from server rack to extract heat from the system

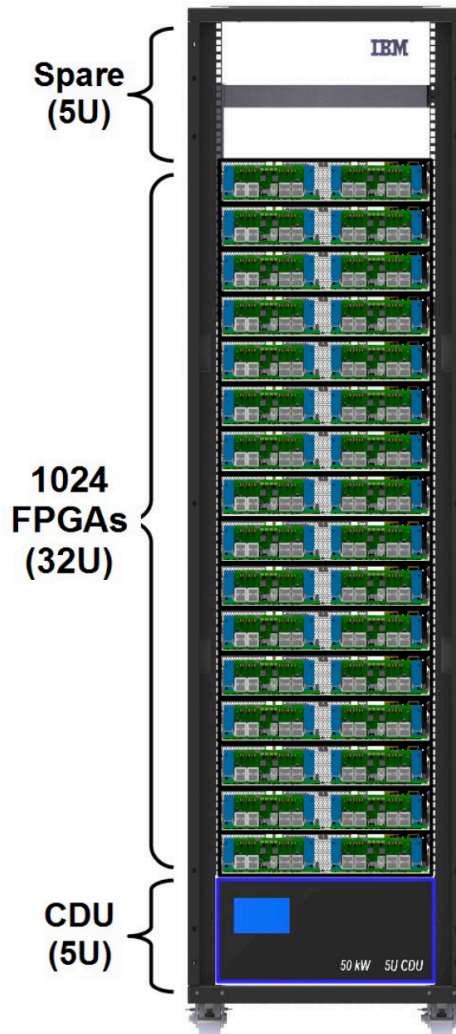



Share:  
[bit.ly/cloudfpga](https://bit.ly/cloudfpga)

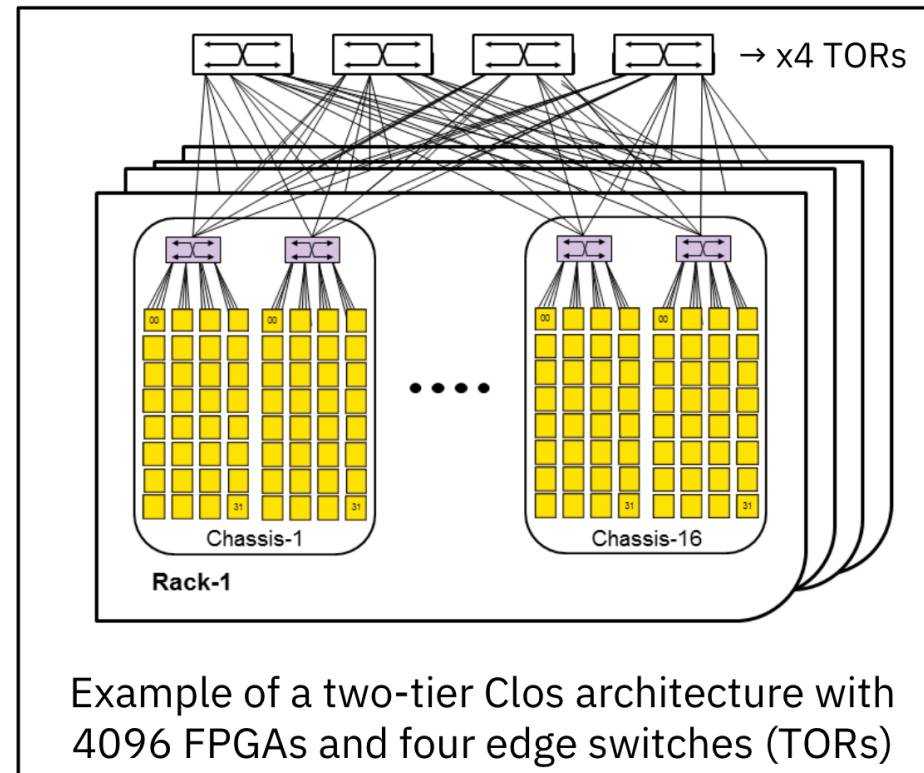
EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019

Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>

# Towards 1024 FPGAs per Server Rack



1024 FPGAs → 2.8M DSPs,  
2x10<sup>15</sup> Fixed-Point Multiply-Accumulates/s  
10 Tb/s bi-sec. Bw – 16 TB DDR4 – 40 kW max.




Share:  
[bit.ly/cloudfpga](https://bit.ly/cloudfpga)

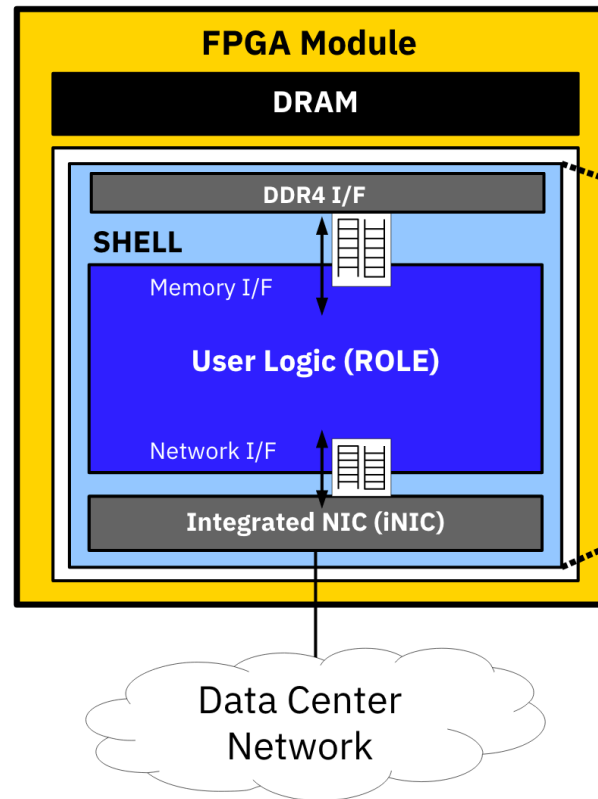
EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019

Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>

# FPGA as a Service using Disaggregated FPGAs




- The FPGAs from the server rack can be assigned to users to create FPGA as a Service setup
- User programs the FPGA with their desired logic




**SHELL** → abstracts the HW components of the FPGA and exposes standard AXI(S) interfaces to the user.

**ROLE** → embeds user's application logic.



topFMKU60.v




Share:  
[bit.ly/cloudfpga](https://bit.ly/cloudfpga)

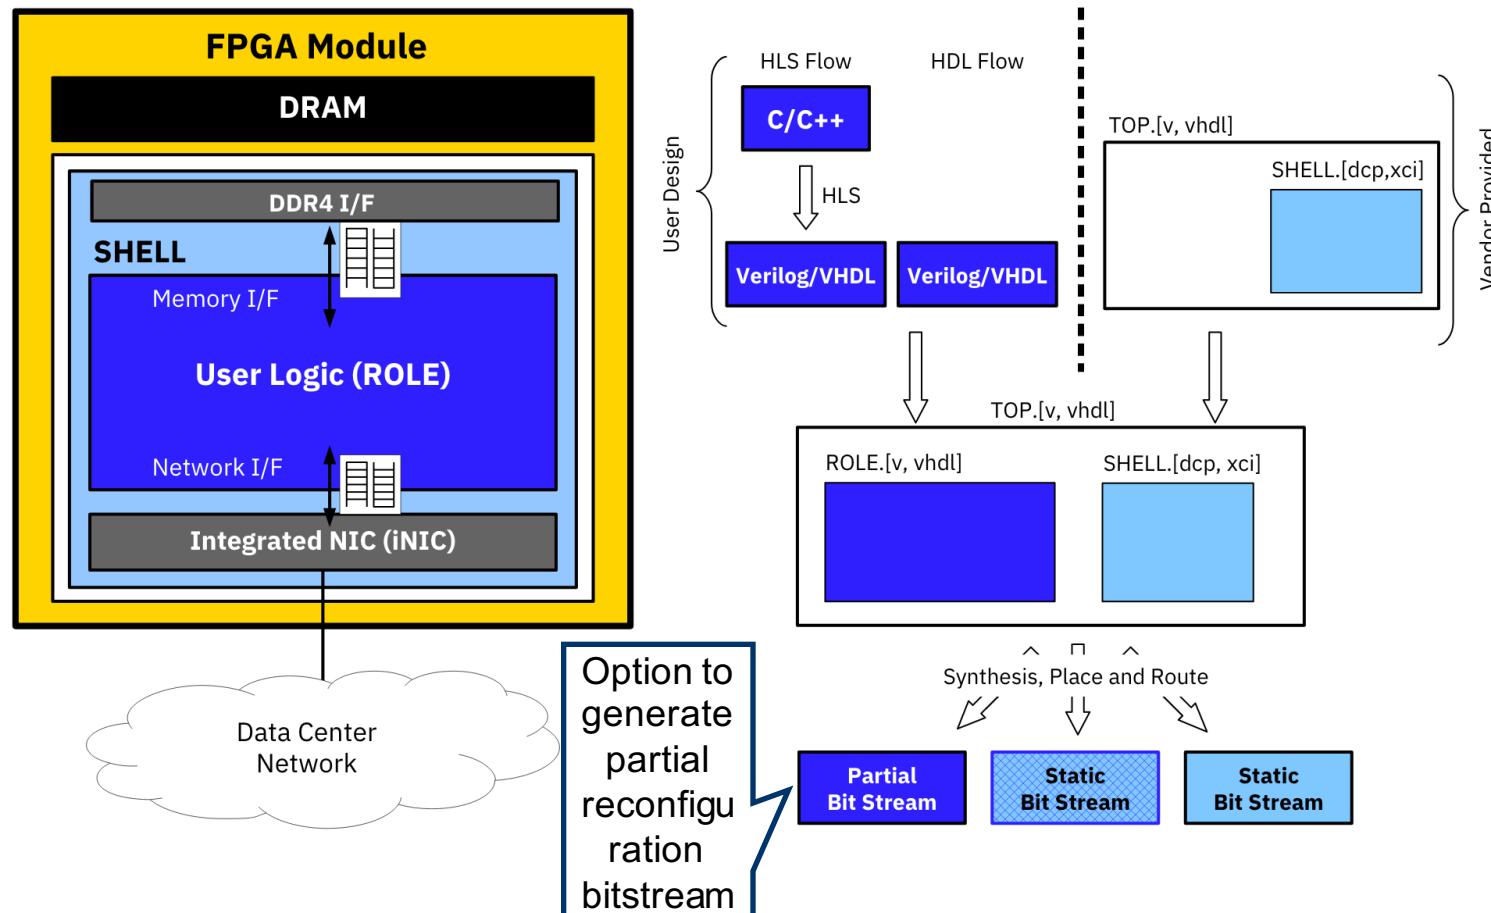
EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019

Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>

# Programming Disaggregated FPGAs



- The FPGAs are programmed in similar manner to other Cloud FPGA deployments

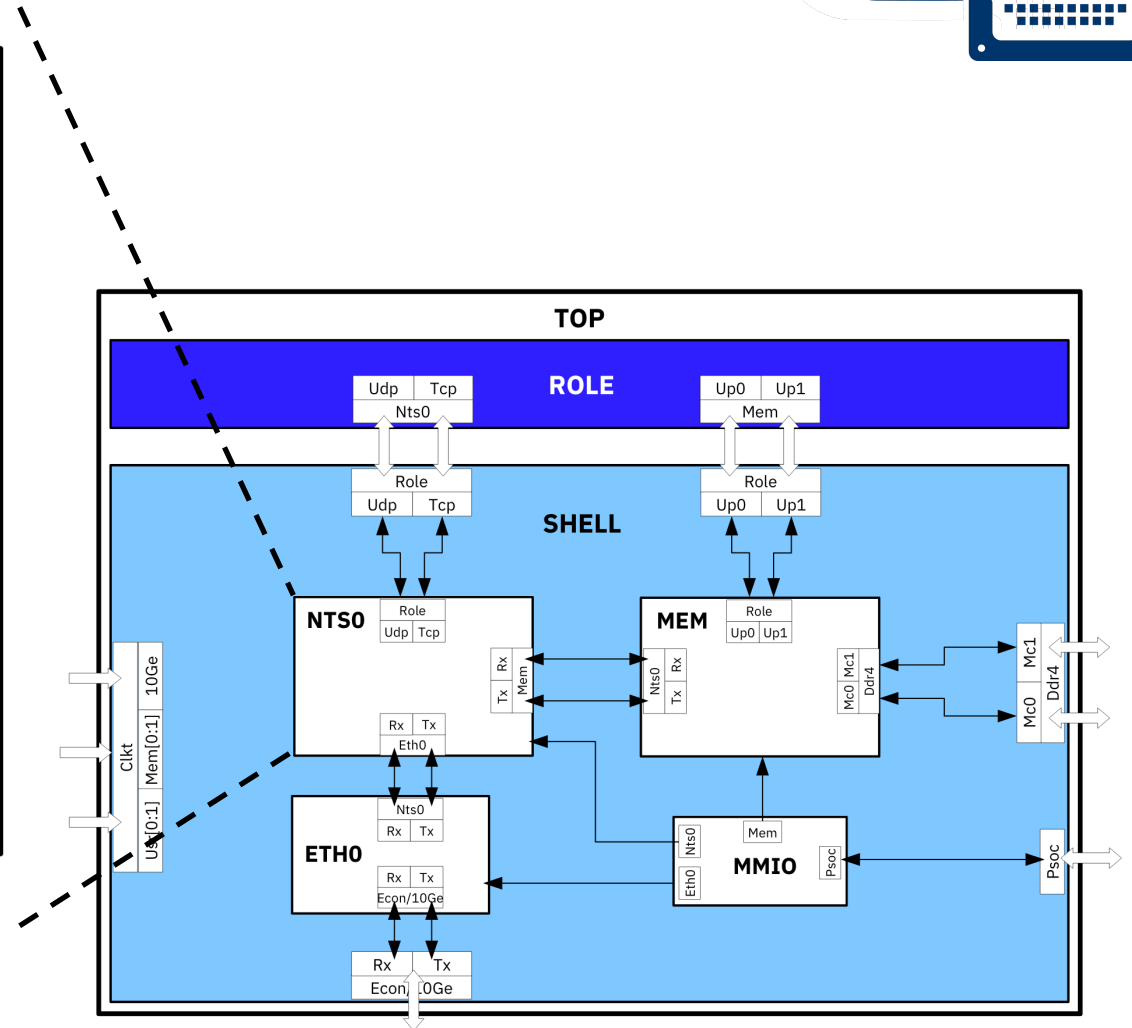
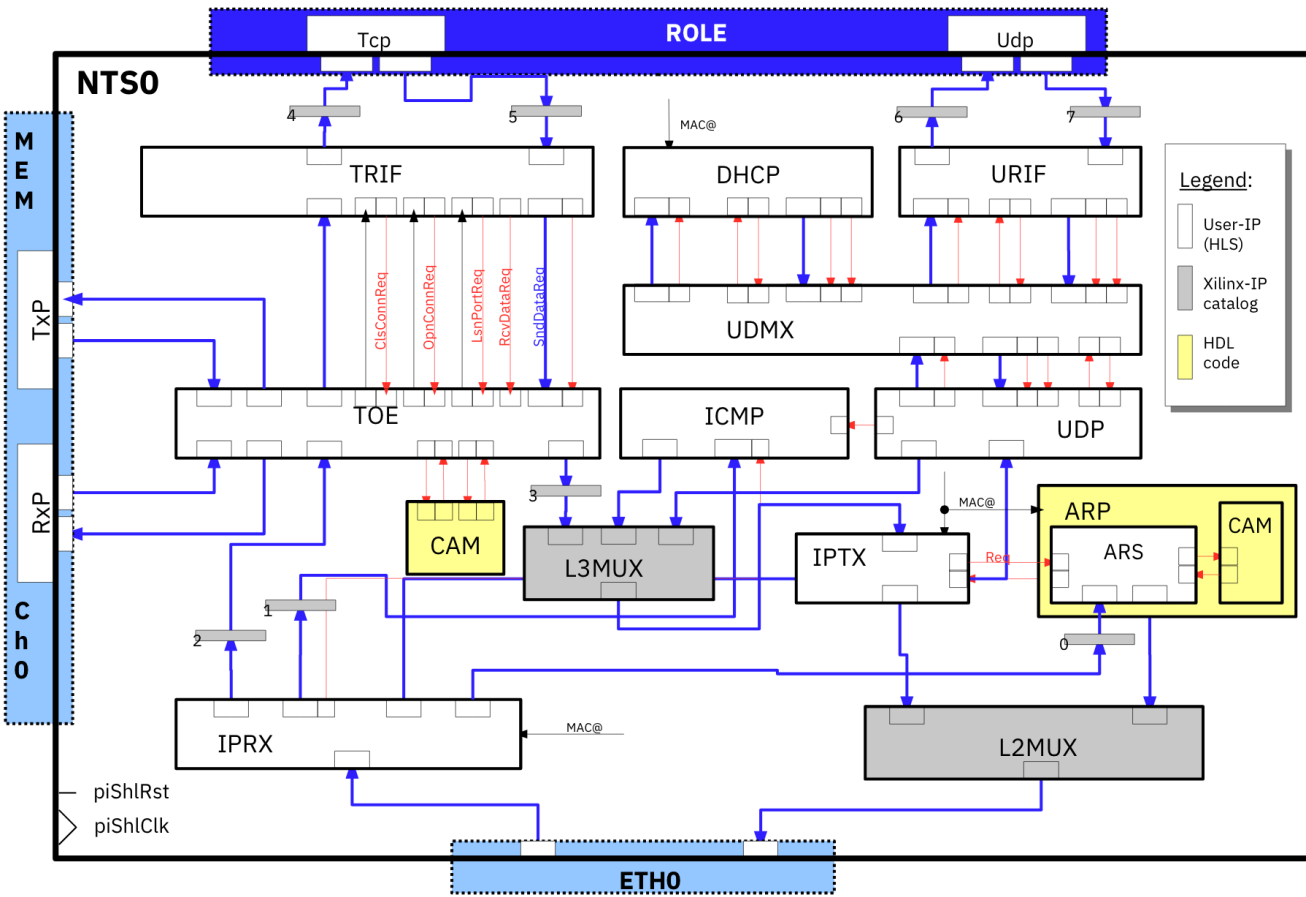



Share:  
[bit.ly/cloudfpga](http://bit.ly/cloudfpga)

EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019


Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>

# Shell Logic and Network-Transport-Session

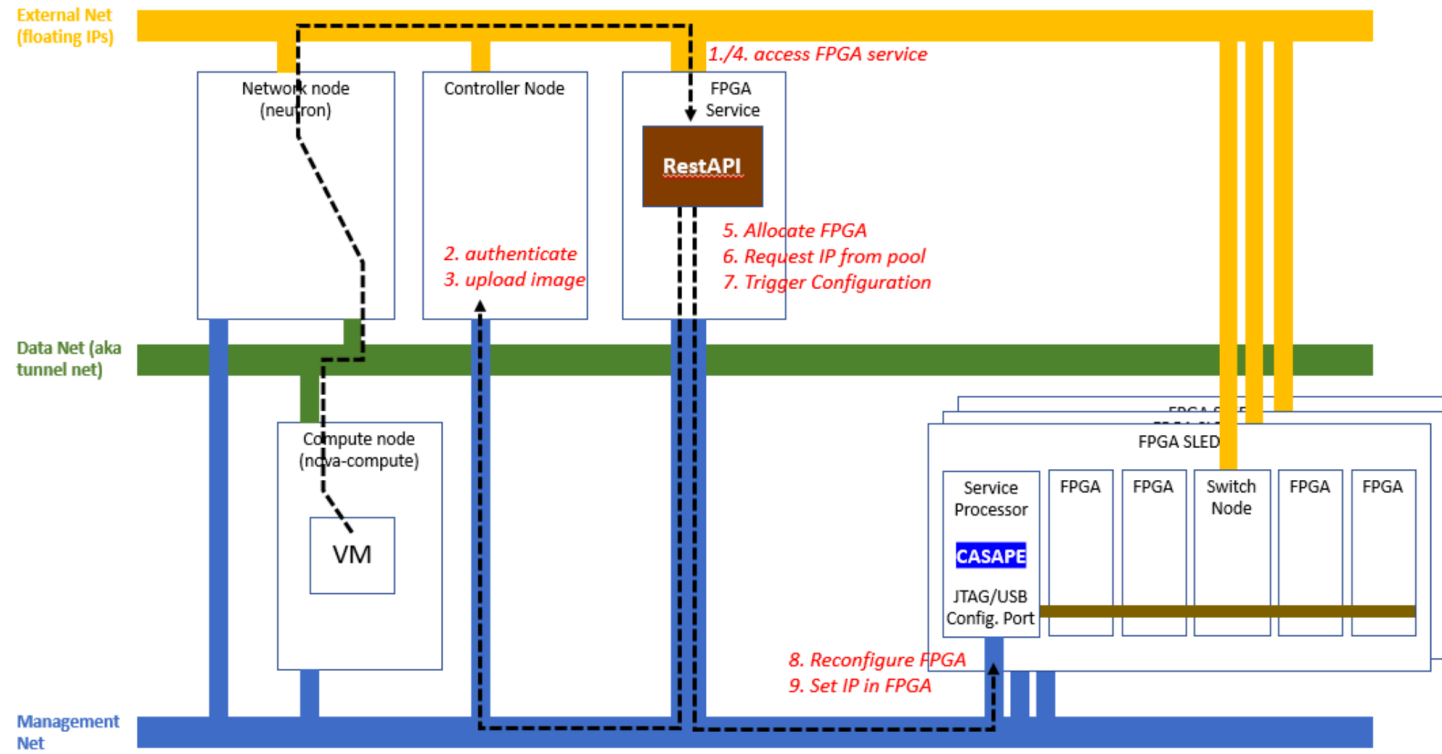


Share: [bit.ly/cloudfpga](https://bit.ly/cloudfpga)

# Networking of the FPGAs



- Separate networks inside the data center are used manage the FPGAs, and to let them communicate




Share:  
[bit.ly/cloudfpga](https://bit.ly/cloudfpga)

EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019

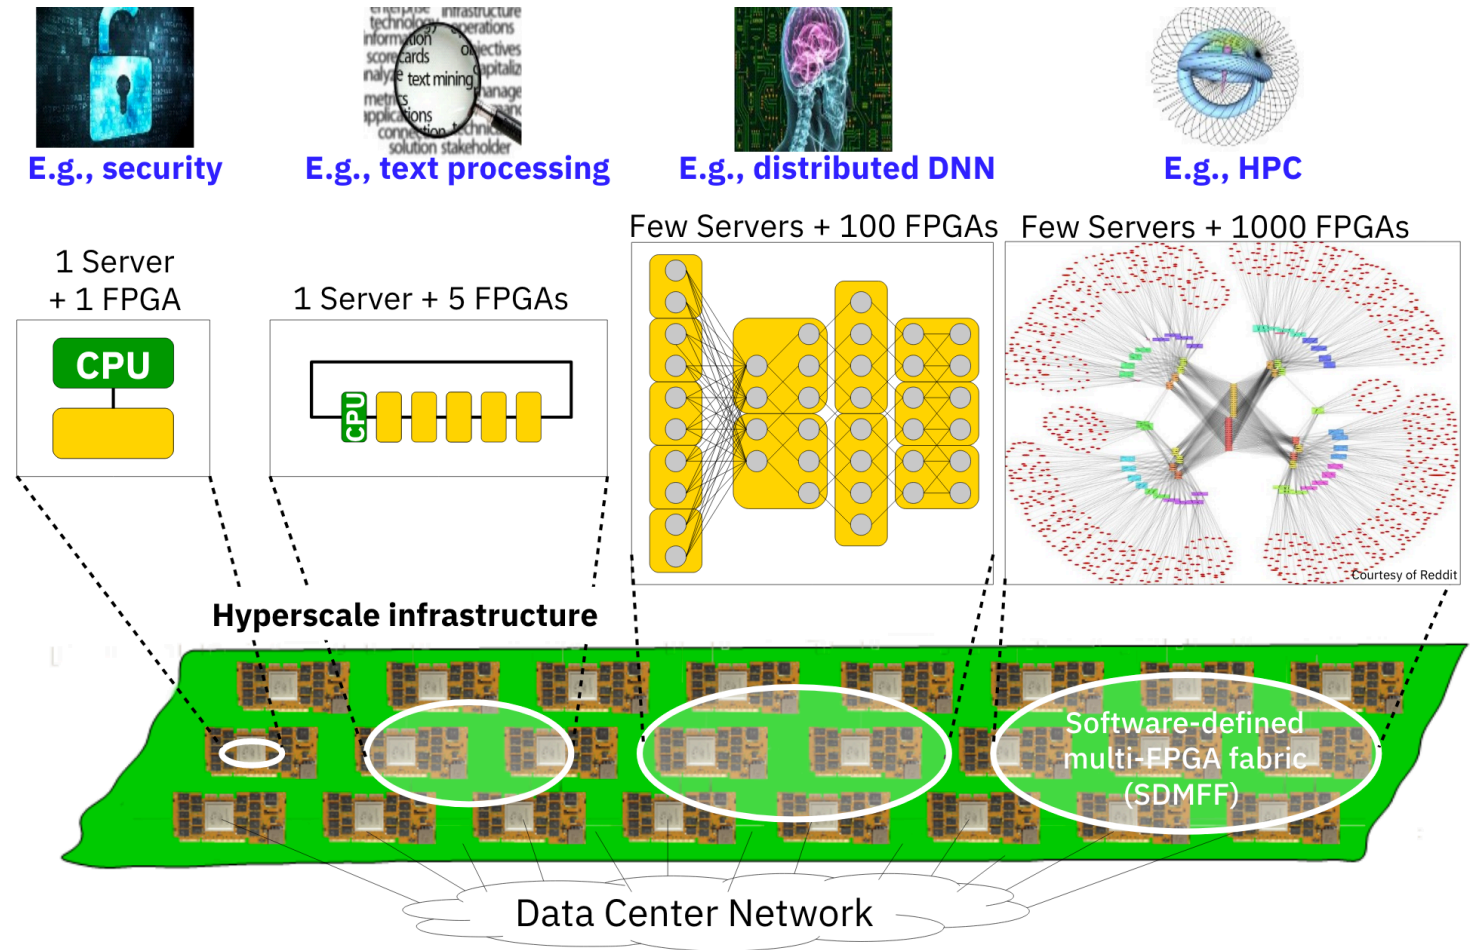
Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>



# Towards Software-Defined Multi-FPGA Fabric



- Disaggregated FPGAs allow for arbitrary connection of the FPGAs to achieve different tasks




Share:  
[bit.ly/cloudfpga](https://bit.ly/cloudfpga)

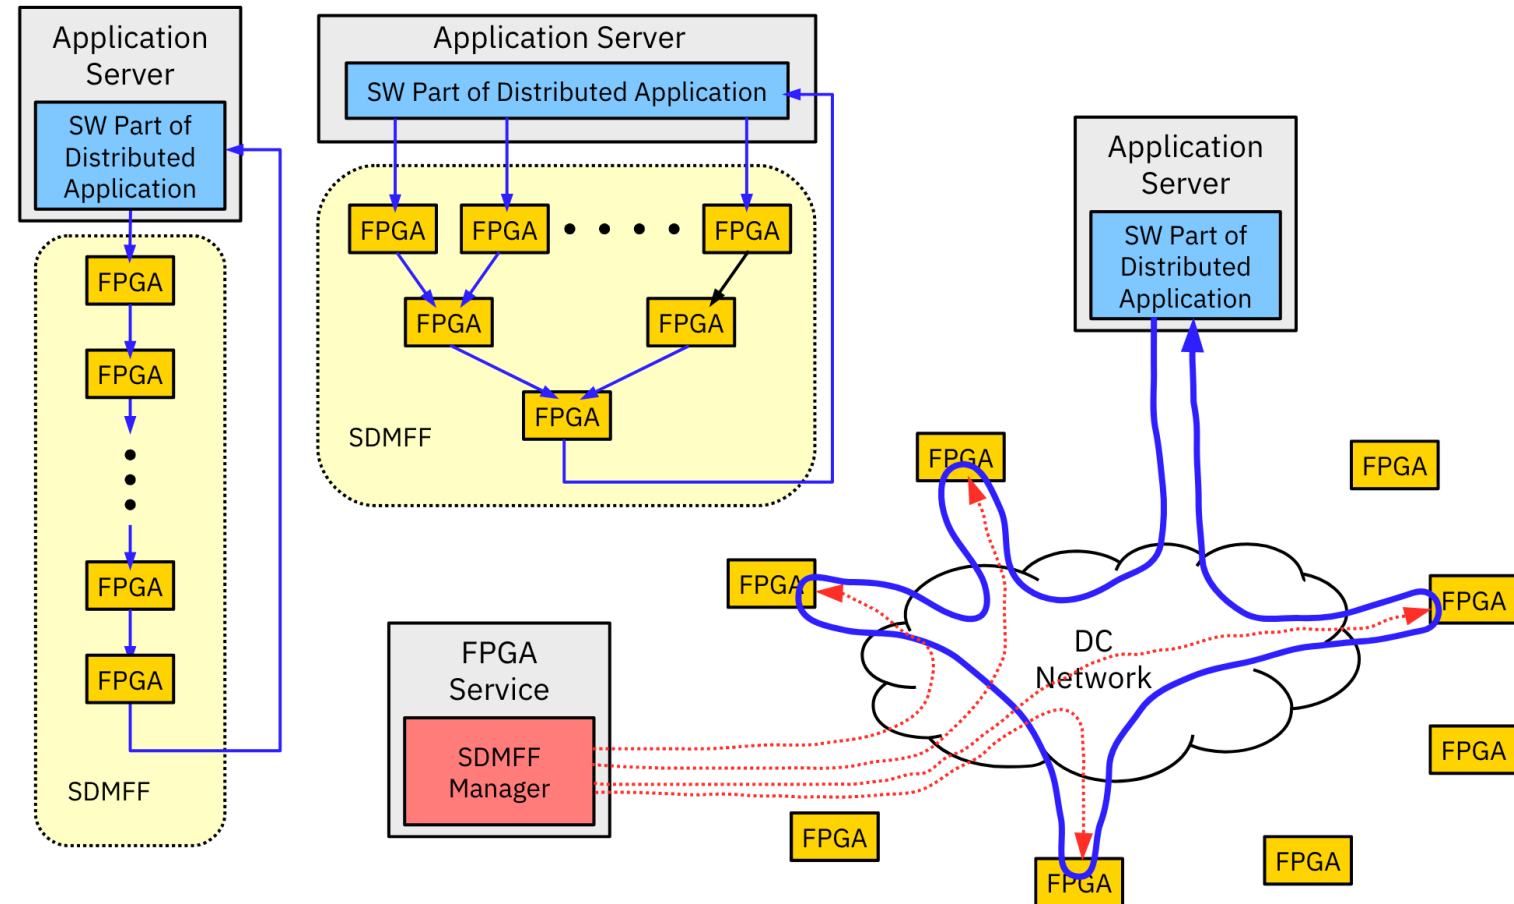
EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019

Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>

# Towards Software-Defined Multi-FPGA Fabric




- FPGA and CPU interaction can be flexibly defined
  - No longer is it CPU controlling what FPGA does



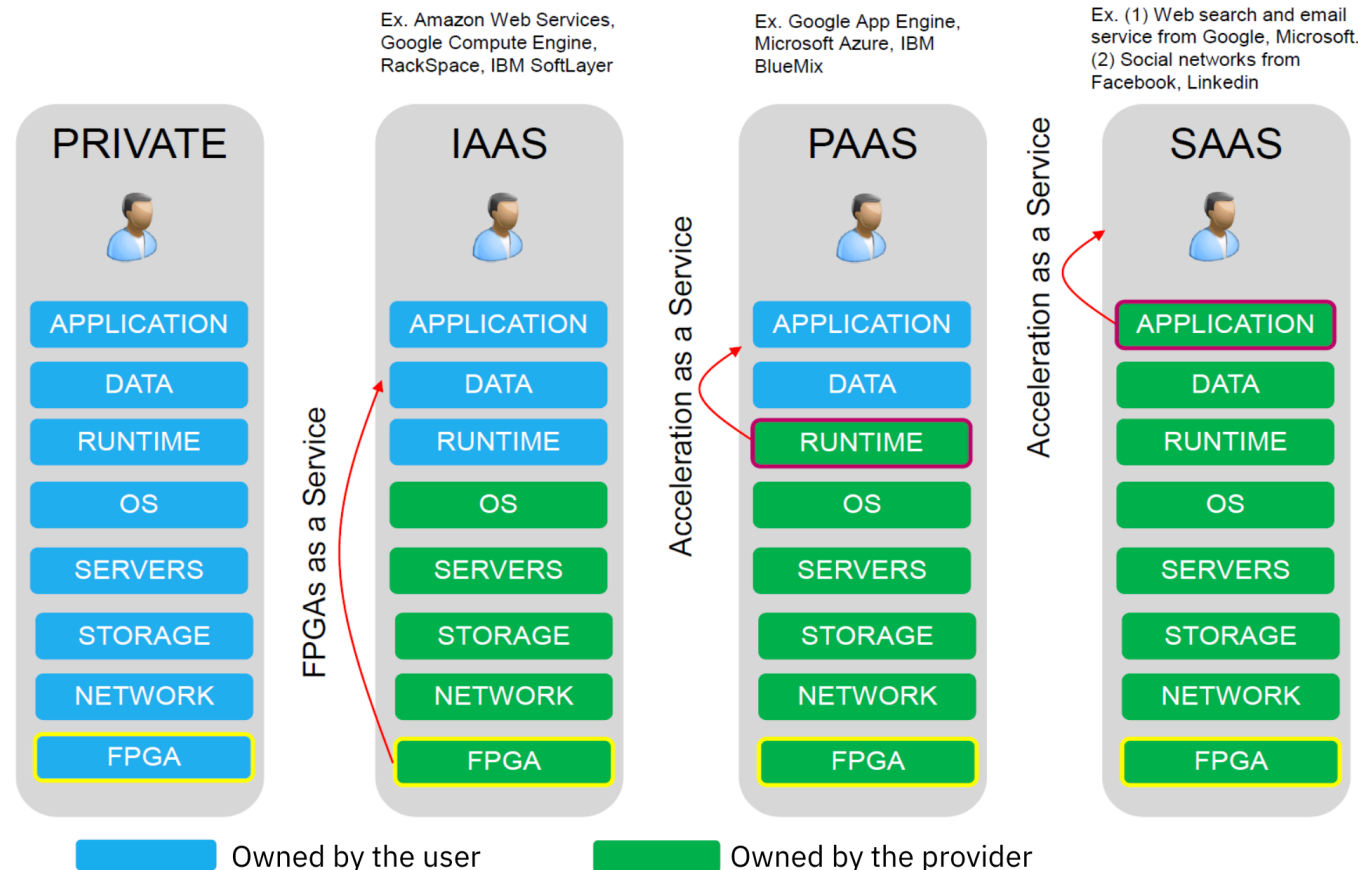
Share:  
[bit.ly/cloudfpga](http://bit.ly/cloudfpga)



# From FPGA as a Service to Acceleration as a Service



- FPGAs in data center can fit into the typical IaaS, PaaS, SaaS models:




Share:  
[bit.ly/cloudfpga](http://bit.ly/cloudfpga)

EENG 428 / ENAS 968 – Cloud FPGA  
© Jakub Szefer, Fall 2019

Based on slides provided courtesy of Francois Abel and Christoph Hagleitner, IBM Zurich, original slides available at:  
<https://www.researchgate.net/project/cloudFPGA>

# References



1. “Data Warehouse Appliance”. Available at: [https://en.wikipedia.org/wiki/Data\\_warehouse\\_appliance](https://en.wikipedia.org/wiki/Data_warehouse_appliance)
2. “IBM Netezza 100 model summary”. Available at: [https://www.ibm.com/support/knowledgecenter/en/SSULQD\\_7.1.0/com.ibm.nz.gsg.doc/c\\_getstrt\\_nz\\_100\\_summary.html](https://www.ibm.com/support/knowledgecenter/en/SSULQD_7.1.0/com.ibm.nz.gsg.doc/c_getstrt_nz_100_summary.html)
3. “IBM PureData System for Analytics N1001 model summary”. Available at: [https://www.ibm.com/support/knowledgecenter/en/SSULQD\\_7.1.0/com.ibm.nz.gsg.doc/c\\_getstrt\\_nz\\_1000\\_summary.html](https://www.ibm.com/support/knowledgecenter/en/SSULQD_7.1.0/com.ibm.nz.gsg.doc/c_getstrt_nz_1000_summary.html)
4. “Netezza”. Available at: <https://en.wikipedia.org/wiki/Netezza>
5. The Register. “Netezza slims TwinFin analytics appliance”. Available at: [https://www.theregister.co.uk/2010/01/27/netezza\\_skimmer\\_appliance/](https://www.theregister.co.uk/2010/01/27/netezza_skimmer_appliance/)
6. “IBM PureData System for Analytics Architecture”. Available at: <https://www.redbooks.ibm.com/redpapers/pdfs/redp4725.pdf>

