Wendelin Exanalytics

Hypercube Big Data Center

2014-06-16 – Shanghai

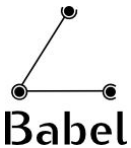

www.wendelin.io

Standard Bay

10 Gbps L2 Switch

Standard Server

IPMI L2 Switch

Management Server

Standard Bay

WENDELIN

Standard Server

- **Minimum**
  - **2 x 10 Gbps**
  - **1 x IPMI**
  - **1 x i7 CPU**
  - **32 GB RAM**
  - **1 x 200 GB SSD**
- **Typical**
  - **2 x 10 Gbps**
  - **1 x IPMI**
  - **2 x 6 core Xeon CPU**
  - **512 GB RAM**
  - **4 x 1 TB SSD**
  - **1 x M2090 GPU**

CORETO

lenovo 联想

inspur 浪潮

WENDELIN

Standard Server

100% open source

- **GNU/Linux** base OS
- **SlapOS** cloud / orchestrator / billing
- **Babel** low latency routing
- **re6st** address range allocation
- **IPMI** management protocol

WENDELIN

Management Server

- **Minimum**
  - **2 x 100 Mbps**
  - **1 x IPMI**
  - **1 x Atom CPU**
  - **1 GB RAM**
  - **1 x 16 GB SSD**
- **Typical**
  - **2 x 1 Gbps**
  - **1 x IPMI**
  - **1 x i5 CPU**
  - **16 GB RAM**
  - **1 x 120 GB SSD**

WENDELIN

Management Server

100% open source

- **GNU/Linux** base OS + router (dhcpd, PXE, iptables)
- **Kadeploy** base OS deployment
- **IPMI** management protocol

WENDELIN

LAN Cabling

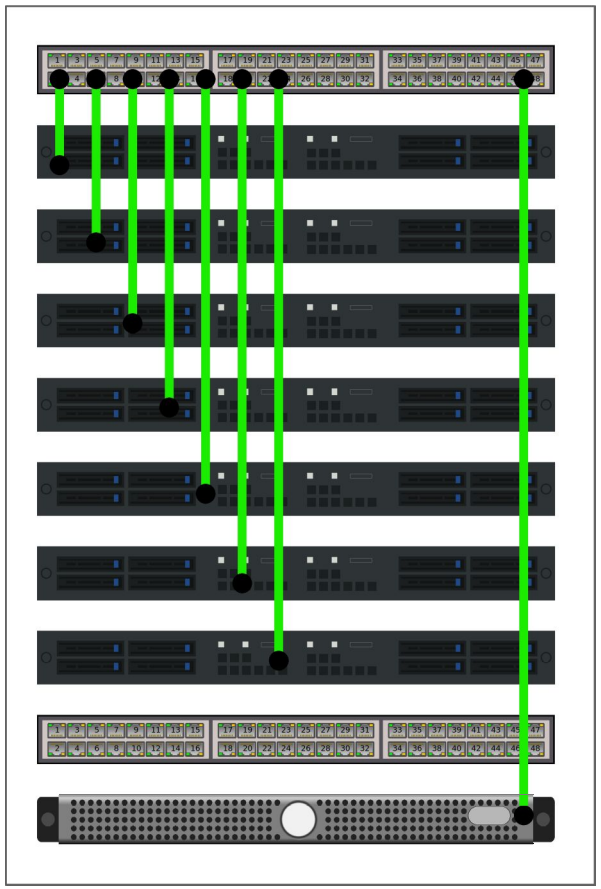

10 Gbps L2 Switch

Standard Server

IPMI L2 Switch

Management Server

Standard Bay

WENDELIN

LAN Configuration

- **Static IPv6 address range**
  - **Big Data transfer**
- **Non routable IPv4**
  - **Base OS deployment**

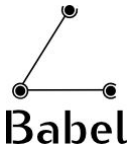

WENDELIN

IPMI Cabling

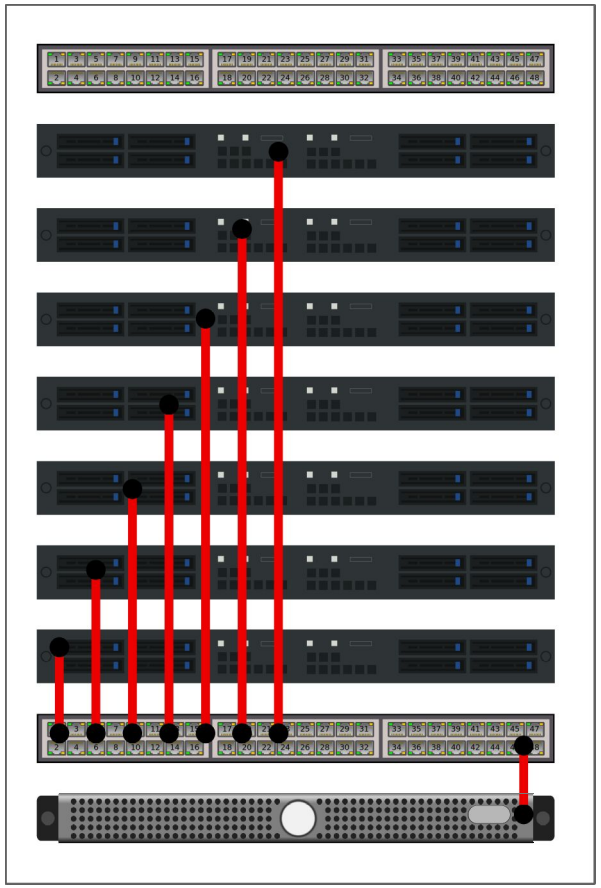

10 Gbps L2 Switch

Standard Server

IPMI L2 Switch

Management Server

Standard Bay

WENDELIN

IPMI Configuration

- **Non routable IPv4**
  - **IPMI access**

WENDELIN

Recursive Management Topology

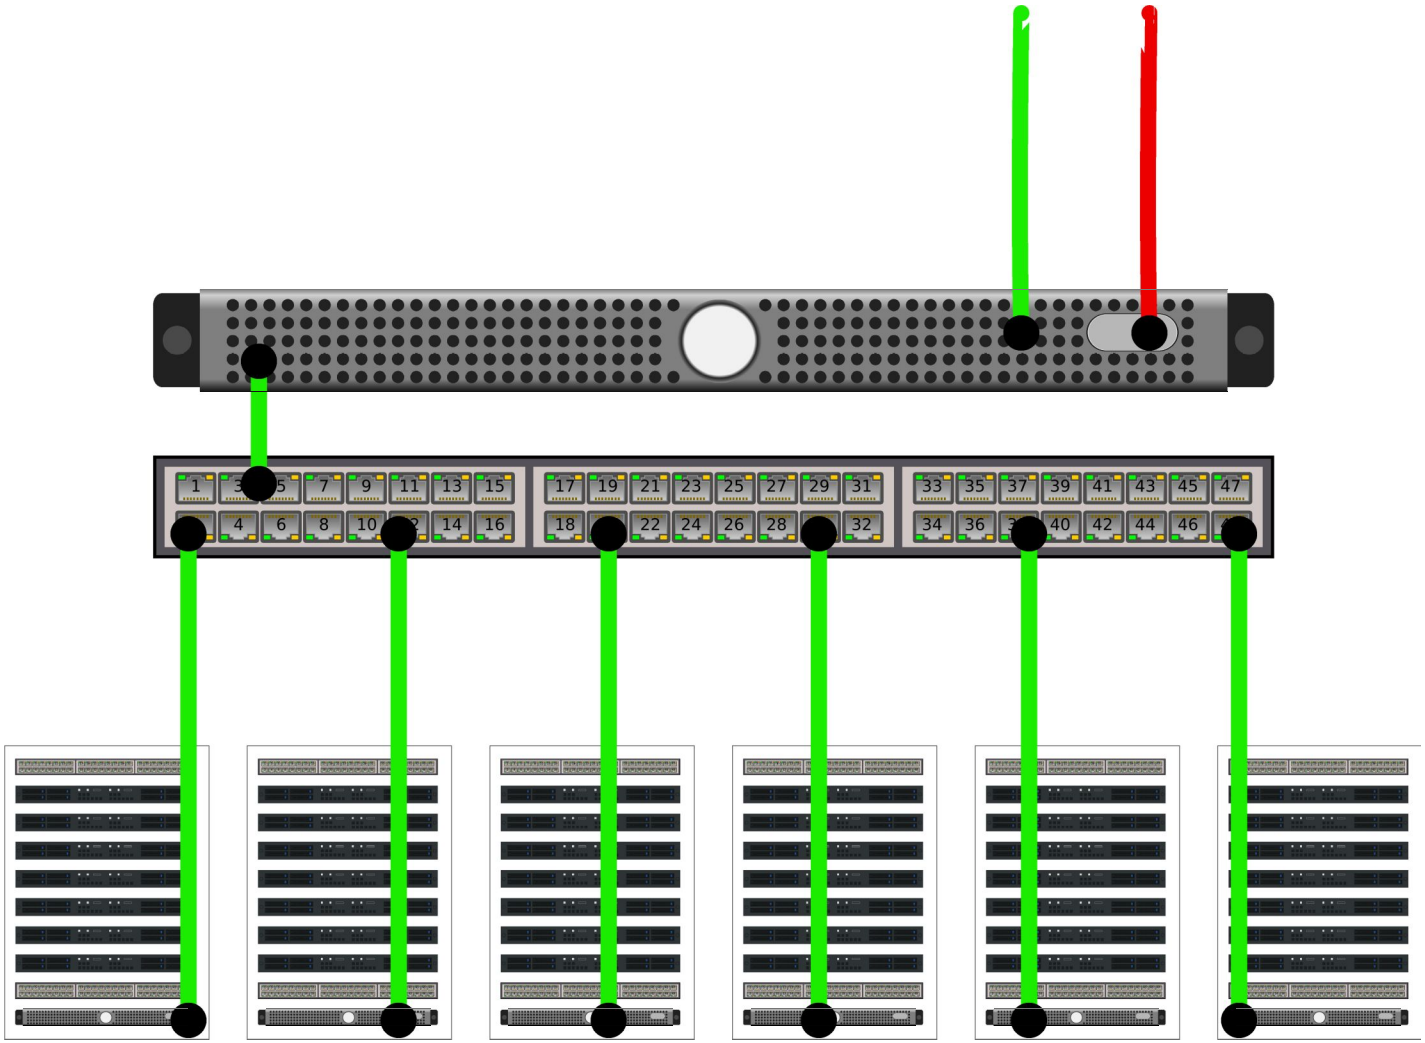

WENDELIN

Hypercube Cabling

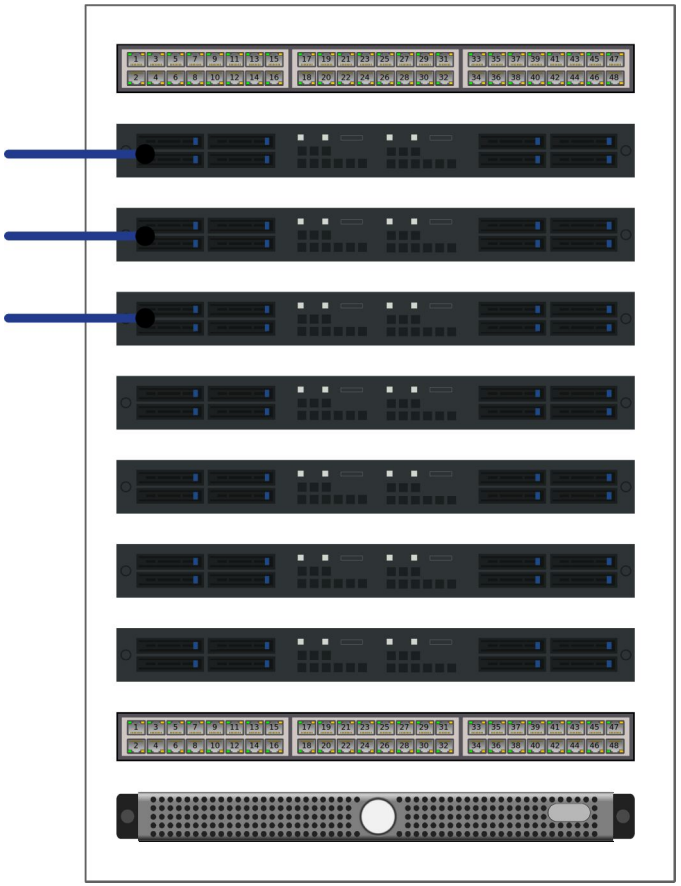

10 Gbps L2 Switch

Standard Server

IPMI L2 Switch

Management Server

Standard Bay

WENDELIN

Hypercube Configuration

- **IPv6 Local Link** no configuration
  - **Big Data transfer**

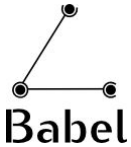

WENDELIN

Hypercube Topology

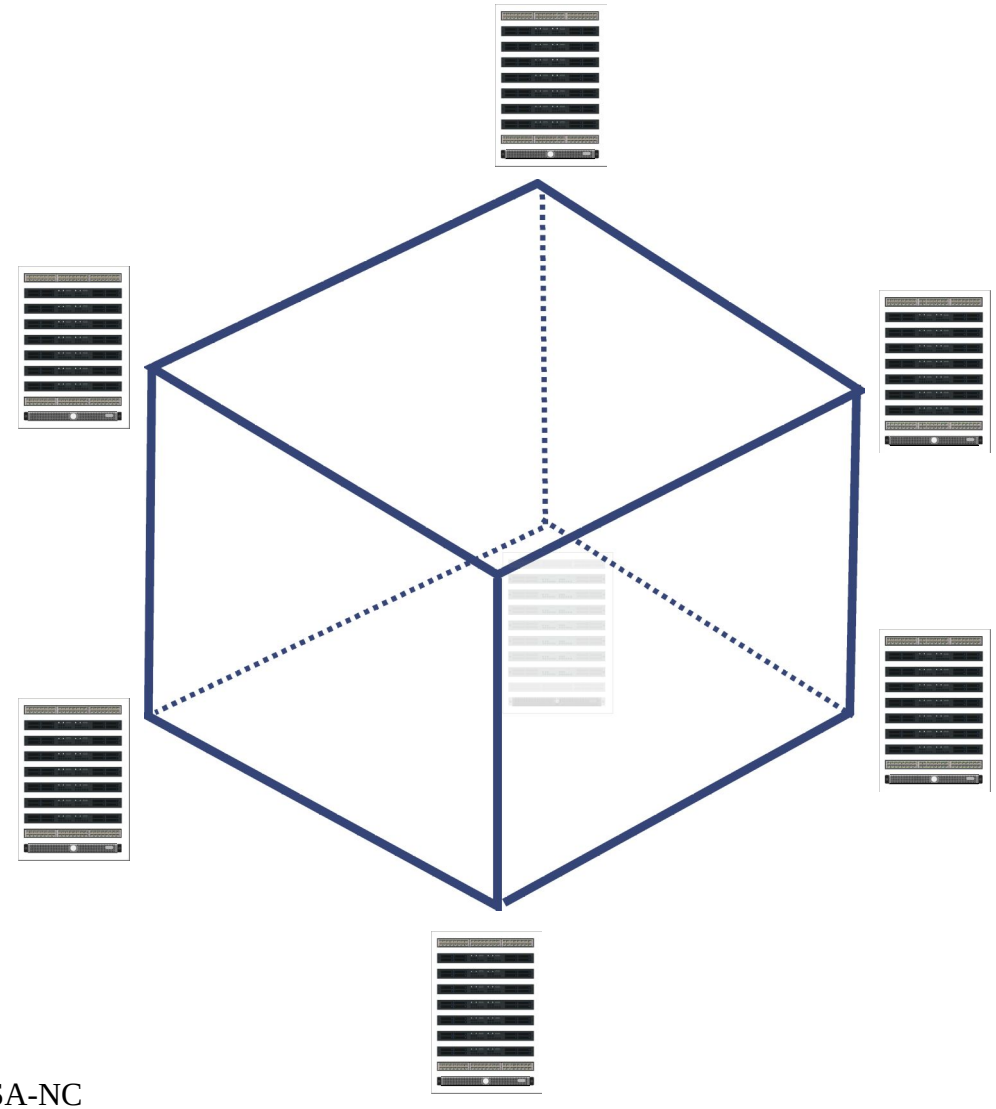

WENDELIN

N-dimensional Cube Performance

- **Size:** $N * 2^N$ exponential
- **Max latency:** $O(N)$ minimal
- **Core bandwidth (Gbps):** $O(2^N)$ scalable
- **Storage Size (TB):** $O(N * 2^N)$ exponential
- **Capacity (GFLOPS):** $O(N * 2^N)$ exponential

WENDELIN

8-dimensional Cube Performance

- **Size:** $N * 2^N$ — 2048 servers
- **Max latency:** $O(N)$ — 8 x (10G + kernel latency)
- **Core bandwidth (Gbps):** $10 * 2^N$ — 2.560 Tbps
- **Storage Size (TB):** $4 * N * 2^N$ — 8192 TB
- **Capacity (GFLOPS):** $500 * N * 2^N$ — 1,024,000 GFLOPS

WENDELIN

Standard Applications

- **Wendelin (Big Data as a Service)**
- **Kvm (Infrastructure as as Service)**
- **Webrunner (Platform as a Service)**
- **MariaDB / TokuDB (Database)**

WENDELIN

Optional Applications open architecture

- **HADOOP (Big Data as a Service)**
- **OpenStack (Infrastructure as as Service)**
- **Mongo (Database)**
- **Cassandra (Database)**
- **PostgreSQL (Database)**
- **etc.**

WENDELIN

Wendelin Exanalytics

Hypercube Big Data Center

2014-06-16 – Shanghai

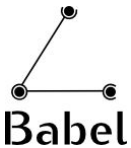

www.wendelin.io

Wendelin Exanalytics

Hypercube Big Data Center

2014-06-16 – Shanghai


www.wendelin.io

Standard Bay

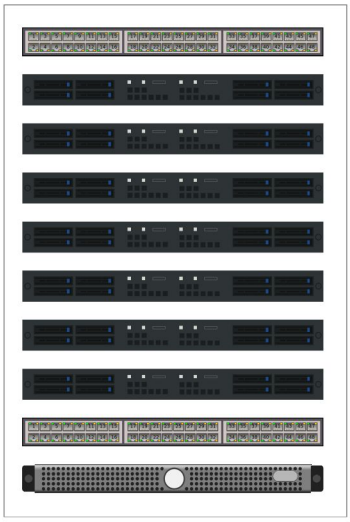

Standard Bay

WENDELIN

Standard Server

- **Minimum**
  - **2 x 10 Gbps**
  - **1 x IPMI**
  - **1 x i7 CPU**
  - **32 GB RAM**
  - **1 x 200 GB SSD**
- **Typical**
  - **2 x 10 Gbps**
  - **1 x IPMI**
  - **2 x 6 core Xeon CPU**
  - **512 GB RAM**
  - **4 x 1 TB SSD**
  - **1 x M2090 GPU**

Standard Server

100% open source

- **GNU/Linux** base OS
- **SlapOS** cloud / orchestrator / billing
- **Babel** low latency routing
- **re6st** address range allocation
- **IPMI** management protocol

Management Server

- **Minimum**
  - **2 x 100 Mbps**
  - **1 x IPMI**
  - **1 x Atom CPU**
  - **1 GB RAM**
  - **1 x 16 GB SSD**
- **Typical**
  - **2 x 1 Gbps**
  - **1 x IPMI**
  - **1 x i5 CPU**
  - **16 GB RAM**
  - **1 x 120 GB SSD**

WENDELIN

Management Server

100% open source

- **GNU/Linux** base OS + router (dhcpd, PXE, iptables)
- **Kadeploy** base OS deployment
- **IPMI** management protocol

LAN Cabling

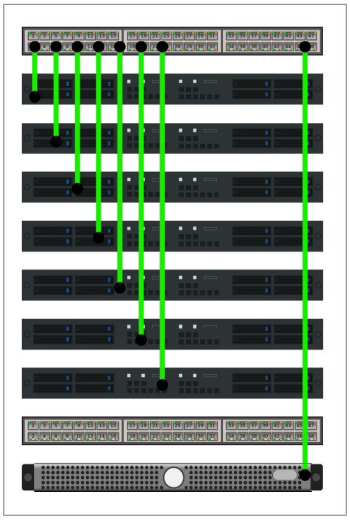

10 Gbps L2 Switch

Standard Server

IPMI L2 Switch

Management Server

Standard Bay

WENDELIN

LAN Configuration

- **Static IPv6 address range** re6st
  - **Big Data transfer**
- **Non routable IPv4** KADEPLOY
  - **Base OS deployment**

WENDELIN

IPMI Cabling

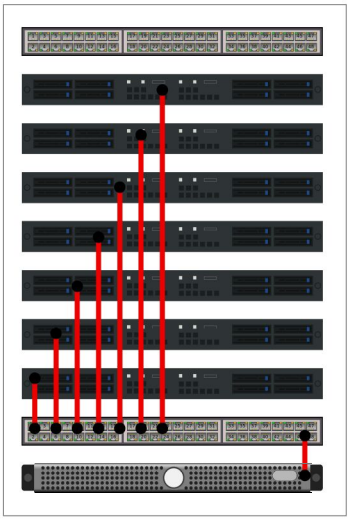

10 Gbps L2 Switch

Standard Server

IPMI L2 Switch

Management Server

Standard Bay

WENDELIN

IPMI Configuration

- **Non routable IPv4**
  - **IPMI access** 

WENDELIN

Recursive Management Topology

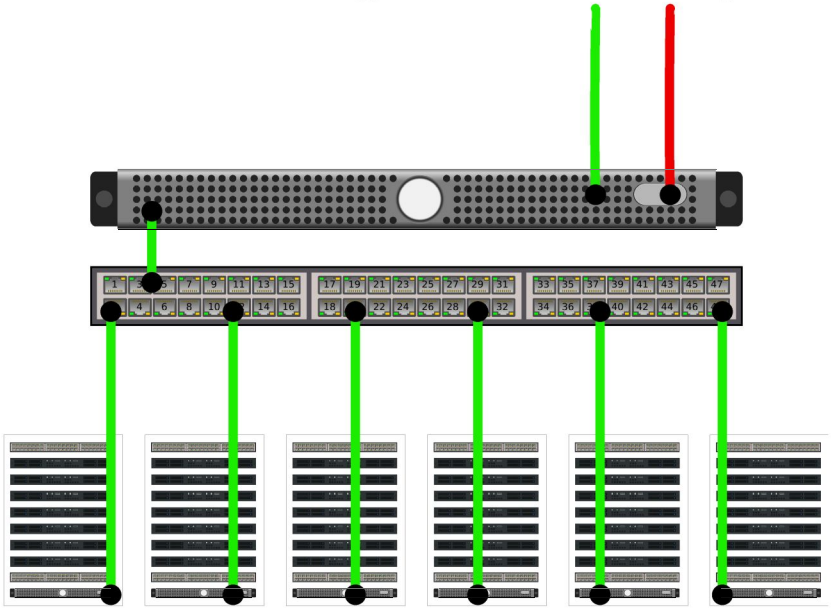

WENDELIN

Hypercube Cabling

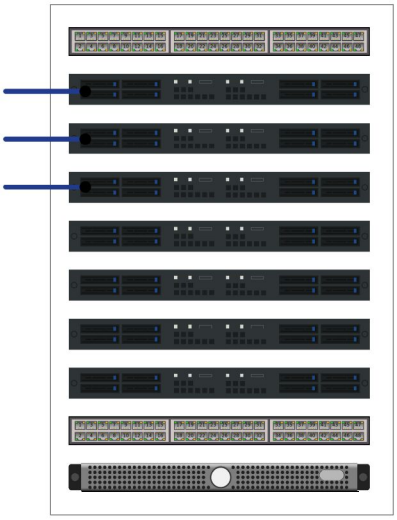

WENDELIN

Hypercube Configuration

- **IPv6 Local Link** no configuration
  - **Big Data transfer**

WENDELIN

Hypercube Topology

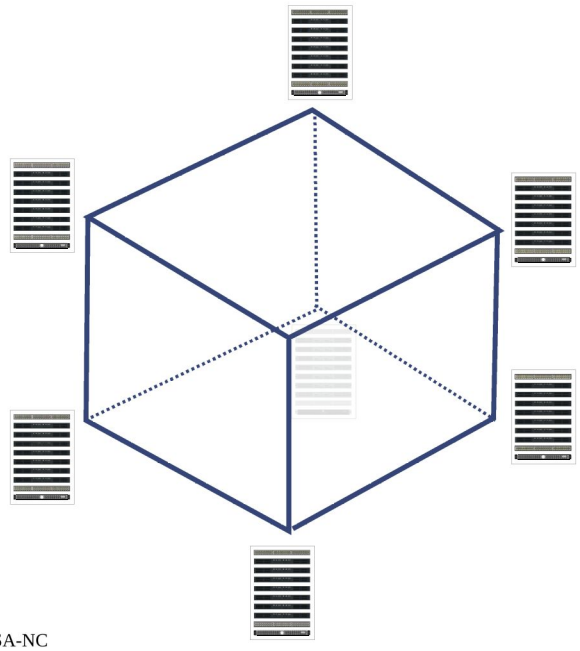

WENDELIN

N-dimensional Cube Performance

- **Size:** $N * 2^N$ exponential
- **Max latency:** $O(N)$ minimal
- **Core bandwidth (Gbps):** $O(2^N)$ scalable
- **Storage Size (TB):** $O(N * 2^N)$ exponential
- **Capacity (GFLOPS):** $O(N * 2^N)$ exponential

8-dimensional Cube Performance

- **Size:** $N * 2^N$ — 2048 servers
- **Max latency:** $O(N)$ — 8 x (10G + kernel latency)
- **Core bandwidth (Gbps):** $10 * 2^N$ — 2.560 Tbps
- **Storage Size (TB):** $4 * N * 2^N$ — 8192 TB
- **Capacity (GFLOPS):** $500 * N * 2^N$ — 1,024,000 GFLOPS

WENDELIN

Standard Applications

- **Wendelin (Big Data as a Service)**
- **Kvm (Infrastructure as as Service)**
- **Webrunner (Platform as a Service)**
- **MariaDB / TokuDB (Database)**

Optional Applications open architecture

- **HADOOP (Big Data as a Service)**
- **OpenStack (Infrastructure as as Service)**
- **Mongo (Database)**
- **Cassandra (Database)**
- **PostgreSQL (Database)**
- **etc.**

WENDELIN

Wendelin Exanalytics

Hypercube Big Data Center

2014-06-16 – Shanghai


www.wendelin.io