Deploying a DCC Infrastructure:
Who and at What Cost?

Adrien Lebre Anthony Simonet Anne-Cécile Orgerie

Inria, Mines Nantes, LINA

Intercloud 2016, Berlin – April 4th, 2016

Introduction: Distributed Cloud Infrastructures

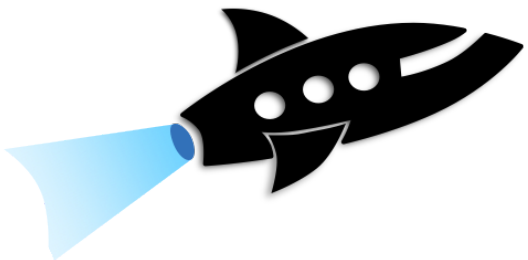

The **Discovery** Initiative: deploying a **fully distributed IaaS infrastructure** on top of network Points of Presence, at the edge of the Internet.

Partners: *Orange, Renater*

From this:

credits: datacentertalk.com - Microsoft DC, Quincy, WA state

To that:

credits: Orange - Orange network PoP

Introduction: Distributed Cloud Infrastructures

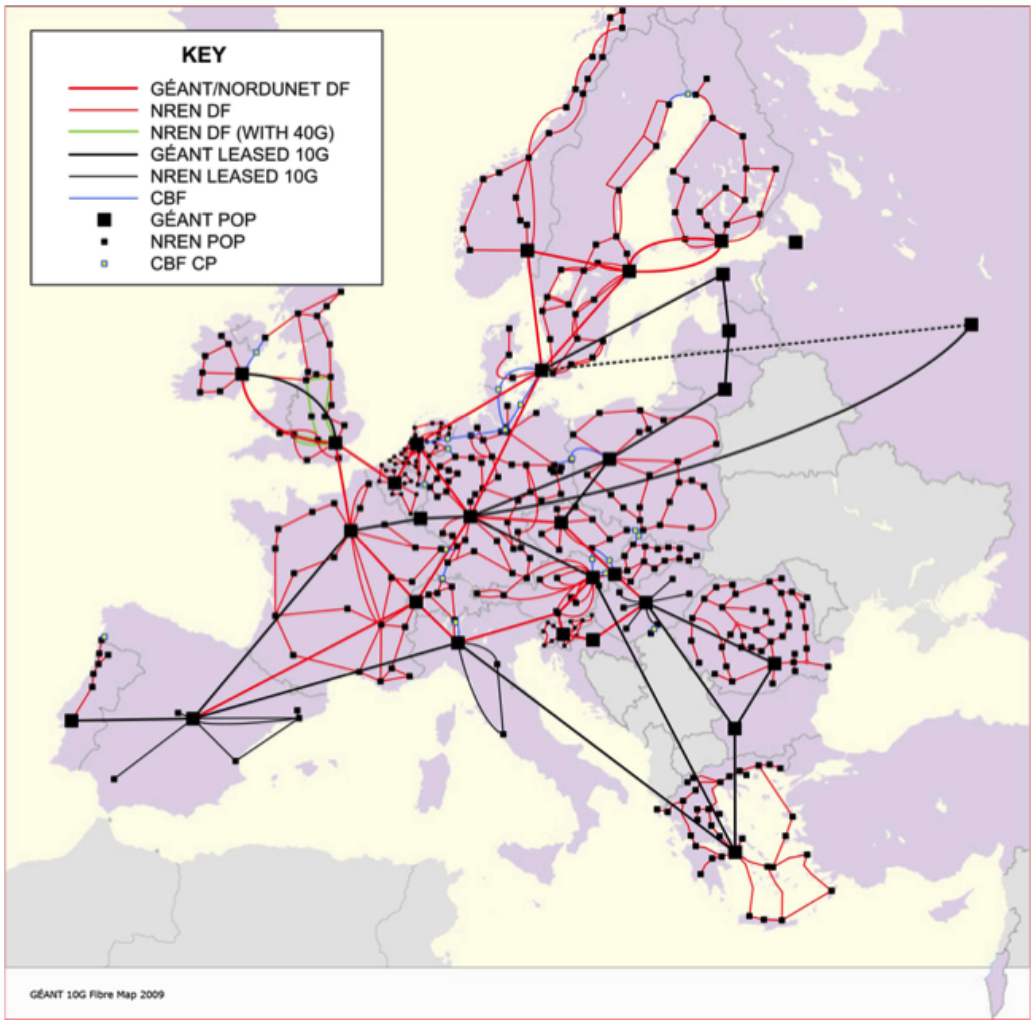

Pros:

- Closer to users (legislation and latency friendly)
- DCs easier to deploy
- Reduces usage of backbone network
- Favor use of local renewable energy

Cons:

- More difficult to operate
- Security is challenging
- What about energy efficiency?

Figure 1: The Géant network

Introduction: Why such a study?

Typical Distributed Cloud questions:

- Will it work?
- Who would/could do that?
- What would the infrastructure look like?
- Would it be economically viable?

Introduction: Why such a study?

Typical Distributed Cloud questions:

- Will it work?
- **Who would/could do that?**
- **What would the infrastructure look like?**
- **Would it be economically viable?**

Introduction: Why such a study?

Typical Distributed Cloud questions:

- Will it work?
- **Who would/could do that?**
  - Cloud Computing Providers
  - Internet Service Providers
  - Telecom Operators
  - New players
- **What would the infrastructure look like?**
- **Would it be economically viable?**

Introduction: Why such a study?

Typical Distributed Cloud questions:

- Will it work?
- **Who would/could do that?**
- **What would the infrastructure look like?**
- **Would it be economically viable?**

Contributions

- We identify the **possible actors** for deploying a DCC infrastructure
- We develop a **cost model** for DCC infrastructures
- We **compare** the cost of 3 scenarios with an industry leader (Amazon EC2)

Classification of actors: What do they need?

A DCC infrastructure requires 3 things:

- Buildings
- Computing, storage and network resources to put in them
- A network to interconnect them

Each of these elements can be operated by actors, or outsourced.

The **cost** and **properties** of the infrastructure will depend on the **level of control** actors have on these elements.

Classification of actors: Level of control

Level of control	Outsourced → Controlled				
	1	2	3	4	5
Inter-DC network	✗	✗	✓	✗	✓
Compute/storage/ intra-DC network	✗	✓	✓	✓	✓
Building	✗	✗	✗	✓	✓

Table 1: Level of control on elements of the infrastructure, from fully outsourced to fully controlled.

Classification of actors: Level of control

Level of control	Outsourced → Controlled				
	1	2	3	4	5
Inter-DC network	✗	✗	✓	✗	✓
Compute/storage/ intra-DC network	✗	✓	✓	✓	✓
Building	✗	✗	✗	✓	✓

Table 1: Level of control on elements of the infrastructure, from fully outsourced to fully controlled.

Classification of actors: Level of control

Level of control	Outsourced → Controlled				
	1	2	3	4	5
Inter-DC network	✗	✗	✓	✗	✓
Compute/storage/ intra-DC network	✗	✓	✓	✓	✓
Building	✗	✗	✗	✓	✓

Table 1: Level of control on elements of the infrastructure, from fully outsourced to fully controlled.

Classification of actors: Level of control

Level of control	Outsourced → Controlled				
	1	2	3	4	5
Inter-DC network	✗	✗	✓	✗	✓
Compute/storage/ intra-DC network	✗	✓	✓	✓	✓
Building	✗	✗	✗	✓	✓

Table 1: Level of control on elements of the infrastructure, from fully outsourced to fully controlled.

Classification of actors: Level of control

Level of control	Outsourced → Controlled				
	1	2	3	4	5
Inter-DC network	✗	✗	✓	✗	✓
Compute/storage/ intra-DC network	✗	✓	✓	✓	✓
Building	✗	✗	✗	✓	✓

Table 1: Level of control on elements of the infrastructure, from fully outsourced to fully controlled.

Classification of actors: Level of control

Level of control	Outsourced → Controlled				
	1	2	3	4	5
Inter-DC network	✗	✗	✓	✗	✓
Compute/storage/ intra-DC network	✗	✓	✓	✓	✓
Building	✗	✗	✗	✓	✓

Table 1: Level of control on elements of the infrastructure, from fully outsourced to fully controlled.

Classification of actors: Qualitative metrics

Level of control					
	1	2	3	4	5
Use-case	New market actor	Deployment over existing facilities, data furnace	Networking operator leveraging existing DC	Deployment of a new DC	Deployment of a new DC and network
Infrastructure elasticity	+++	+	+	–	– – –
Network supervision (monitoring/-control)	– – –	– –	+++	+	+++

Table 2: Qualitative metrics for actors depending on their level of control on multiple parts of the infrastructure.

Classification of actors: Qualitative metrics

Level of control	Outsourced ⟶ Controlled				
	1	2	3	4	5
Maintenance (Hardware)	Outsourced	Computation, storage, intra-DC network	Computation, storage, intra/inter-DC network	Computation, storage, intra-DC network	Computation, storage, intra/inter-DC network
Power	No control	Limited control	No control	Renewable and choice of provider possible	Renewable and choice of provider possible

Table 2: Qualitative metrics for actors depending on their level of control on multiple parts of the infrastructure.

Classification of actors: Qualitative metrics

Level of control	Outsourced → Controlled				
	1	2	3	4	5
Cooling/PUE	No control	Limited control	Limited control	PUE can benefit from latest achievements (free cooling, etc.)	PUE can benefit from latest achievements (free cooling, etc.)
Security (Human presence)	Outsourced (managed by provider)	Outsourced (managed by provider)	Varying: depends on existing infrastructure and human presence in POP	Costly (need to secure a new POP)	Costly (need to secure a new POP)

Table 2: Qualitative metrics for actors depending on their level of control on multiple parts of the infrastructure.

Classification of actors: Qualitative metrics

Level of control					
	1	2	3	4	5
Network cost	Monthly based	Monthly based	Leveraging existing	Monthly based/outsourced	Initial investment
Reliability/ resiliency	Possibility of several providers	Single network provider	Single provider	Several network providers	Single provider
Coverage ratio	Few locations	Deployed according to the demand	Numerous locations	Deployed according to the demand	Maximum coverage

Table 2: Qualitative metrics for actors depending on their level of control on multiple parts of the infrastructure.

Cost model for DCC: Categories of costs

Cost model for a single DC, with 7 categories of cost:

1. Servers
2. Storage
3. Network
4. Power
5. Cooling
6. Building
7. Maintenance

All actors will be subject to these costs.

Cost model for DCC: Categories of costs

Cost model for a single DC, with 7 categories of cost:

1. Servers
2. Storage
3. Network
4. Power
5. Cooling
6. Building
7. Maintenance

All actors will be subject to these costs.

Constant	Type	Usual value
A_s	Servers, routers and switches	5 years
A_m	Racks, cables and CRAC	10 years
A_l	Buildings, backbone network	20 years

Table 3: Amortization constants for different kinds of resources and usual values.

Cost model for DCC: Servers & storage

Servers & storage

We consider 3 kinds of servers:

- Controllers
- Compute servers
- Storage servers

$$Cost_{contr} = N_{contr} \times P_{contr} \times A_m(A_s) \quad (1)$$

$$Cost_{comp} = N_{comp} \times P_{comp} \times A_m(A_s) \quad (2)$$

$$Cost_{storage} = N_{storage} \times P_{storage} \times \underbrace{A_m(A_s)}_{\text{Amortization}} \quad (3)$$

Cost model for DCC: Network

Network

Monthly cost of interconnect **between DCs** and between servers **inside DCs**.
Includes costs of backbone network, routers and switches.

$$Cost_{network} = Cost_{intranet} + Cost_{internet} \tag{4}$$

Equations 5 and 6 detail these two costs:

$$Cost_{intranet} = N_{switch} \times P_{switch} \times A_m(A_s) \tag{5}$$

$$Cost_{internet} = P_{backbone} \times A_m(A_l) \tag{6}$$

Cost model for DCC: Power

Power

Cost of powering the infrastructure, i.e. servers, network devices and racks.
Based on manufacturer data sheets, assuming **100% CPU usage**.

Total **power draw** of the DC in watts:

$$E_{total} = E_{contr} \times N_{contr} + E_{comp} \times N_{comp} + E_{storage} \times N_{storage} + E_{switch} \times N_{switch} + E_{rack} \times N_{rack} \quad (7)$$

Which gives the monthly cost:

$$Cost_{power} = \frac{E_{total}}{1,000} \times 24 \times 30 \times \underbrace{P_{KW}}_{\text{A KW of electricity}} \quad (8)$$

Cost model for DCC: Cooling

Cooling

Add up the cost of **CRACs** and the **electricity** to power them.

1 W of energy consumed produces about **1 W of heat to dissipate.**

$$Cost_{cooling} = \underbrace{N_{CRAC} \times P_{CRAC}}_{\text{Cost of CRACs}} + \underbrace{\frac{E_{total}}{1,000} \times 24 \times 30 \times P_{KW} \times L}_{\text{Cost of power consumed by the CRACs}} \quad (9)$$

The **cooling factor** $L = 0.8$ means 0.8W of power is required to dissipate 1W of heat.

Cost model for DCC: Cooling

Cooling

Add up the cost of **CRACs** and the **electricity** to power them.

1 W of energy consumed produces about **1 W of heat to dissipate.**

$$Cost_{cooling} = N_{CRAC} \times P_{CRAC} + \underbrace{\frac{E_{total}}{1,000} \times 24 \times 30 \times P_{KW} \times L}_{Cost_{power}} \quad (9)$$

The **cooling factor** $L = 0.8$ means 0.8W of power is required to dissipate 1W of heat.

Cost model for DCC: Cooling

Cooling

Add up the cost of **CRACs** and the **electricity** to power them.

1 W of energy consumed produces about **1 W of heat to dissipate.**

$$Cost_{cooling} = N_{CRAC} \times P_{CRAC} + Cost_{power} \times L \quad (9)$$

The **cooling factor** $L = 0.8$ means 0.8W of power is required to dissipate 1W of heat.

Fichera *et al.*, "Power and cooling heat up the data center", 2006

Cost model for DCC: Cooling

Cooling

Add up the cost of **CRACs** and the **electricity** to power them.

1 W of energy consumed produces about **1 W of heat to dissipate.**

$$Cost_{cooling} = N_{CRAC} \times P_{CRAC} + Cost_{power} \times L \tag{9}$$

The **cooling factor** $L = 0.8$ means 0.8W of power is required to dissipate 1W of heat.

$$N_{CRAC} = \lceil E_{total}/1,000/ \underbrace{C_{KW}}_{\substack{\text{CRAC} \\ \text{cooling} \\ \text{capacity}}} \rceil \tag{10}$$

Cost model for DCC: Maintenance

Maintenance

Human cost of **installation** and **maintenance**:

- 1 day of pay per month (on-site interventions)
- 1 full-time engineer for the whole infrastructure

$$Cost_{maintenance} = \underbrace{C_{en}}_{\substack{\text{Monthy} \\ \text{cost of} \\ \text{engineer}}} / \underbrace{N_{DC}}_{\substack{\text{Number} \\ \text{of DCs}}} + \underbrace{C_{tech}}_{\substack{\text{Technician} \\ + \\ \text{travel} \\ \text{expenses}}} \quad (11)$$

Cost model for DCC: Facilities

Facilities

Longest amortization period: 20 years for **buildings**, 10 years for **racks, cables and PDUs**.

$$Cost_{facilities} = \underbrace{P_{building} \times A_m(A_l)}_{\text{Real-estate cost}} + \underbrace{N_{racks} \times P_{rack}}_{\text{Racks}} + \underbrace{N_{switch} \times P_{cables} \times A_m(A_m)}_{\text{Cables}} \quad (12)$$

Use-cases

We instantiate the model for 3 scenarios:

1. a company outsourcing as much as possible (level 1)
2. a NREN (level 3)
3. a Telecom Operator (level 5).

Level of control	 Outsourced → Controlled				
	1	2	3	4	5
Inter-DC network	✗	✗	✓	✗	✓
Compute/storage/ intra-DC network	✗	✓	✓	✓	✓
Building	✗	✗	✗	✓	✓

Use-cases

We instantiate the model for 3 scenarios:

1. a company outsourcing as much as possible (level 1)
2. a NREN (level 3)
3. a Telecom Operator (level 5).

About **10 servers** and **100 VM** per DC. VMs are equivalent or superior to Amazon **m3.medium** instances.

1VCPU, 4GB RAM, 4GB storage

Use-cases: Instantiation

	Facilities	Servers	Storage	Network	Power	Cooling	Maintenance
Outsourced	0€	1,400€	320€	0€	0€	0€	1,000€
Outsourced+BW	0€	1,400€	320€	3,750€	0€	0€	1,000€
NREN	44€	525€	45€	315€	405€	363€	1,500€
Telecom Operator	44€	525€	45€	1354€	405€	363€	1,500€

Table 4: Monthly estimated cost for operating a DC in a DCC infrastructure.

Use-cases: Instantiation

	Facilities	Servers	Storage	Network	Power	Cooling	Maintenance
Outsourced	0€	1,400€	320€	0€	0€	0€	1,000€
Outsourced+BW	0€	1,400€	320€	3,750€	0€	0€	1,000€
NREN	44€	525€	45€	315€	405€	363€	1,500€
Telecom Operator	44€	525€	45€	1354€	405€	363€	1,500€

Table 4: Monthly estimated cost for operating a DC in a DCC infrastructure.

Provider	Cores	One server	1 Gbps	Total	Cost/VM
OVH	8	100€	150€	250€	12.50€
Online	12	150€	N/A	150€	12.50€
Netissime	12	135€	N/A	134€	11.25€
Amen	12	178€	60.20€	238€	14.83€

Table 5: Cost comparison of multiple French providers

Use-cases: Instantiation

	Facilities	Servers	Storage	Network	Power	Cooling	Maintenance
Outsourced	0€	1,400€	320€	0€	0€	0€	1,000€
Outsourced+BW	0€	1,400€	320€	3,750€	0€	0€	1,000€
NREN	44€	525€	45€	315€	405€	363€	1,500€
Telecom Operator	44€	525€	45€	1354€	405€	363€	1,500€

Table 4: Monthly estimated cost for operating a DC in a DCC infrastructure.

	Total	**Cost/VM**
Outsourced	**2,720€**	**26.15€**
Outsourced+BW	**6,470€**	**62.21€**
NREN	**3,197€**	**29.60€**
Telecom Operator	**4,236€**	**39.22€**

Table 5: Total cost and comparison of the cost of a VM with Amazon's 54€.

Conclusion

A few more things to consider:

- All actors on the scale will produce very different infrastructures
- The cost of bandwidth can be mitigated:
  - More BW is more expensive? Charge it to your customers!
  - Amazon **does not** guarantee BW
- **New business models** are to be found

Other aspects of Discovery:

- **Internet-scale P2P Cloud Manager** based on **OpenStack**
- Energy footprint

Thank you!

Questions?

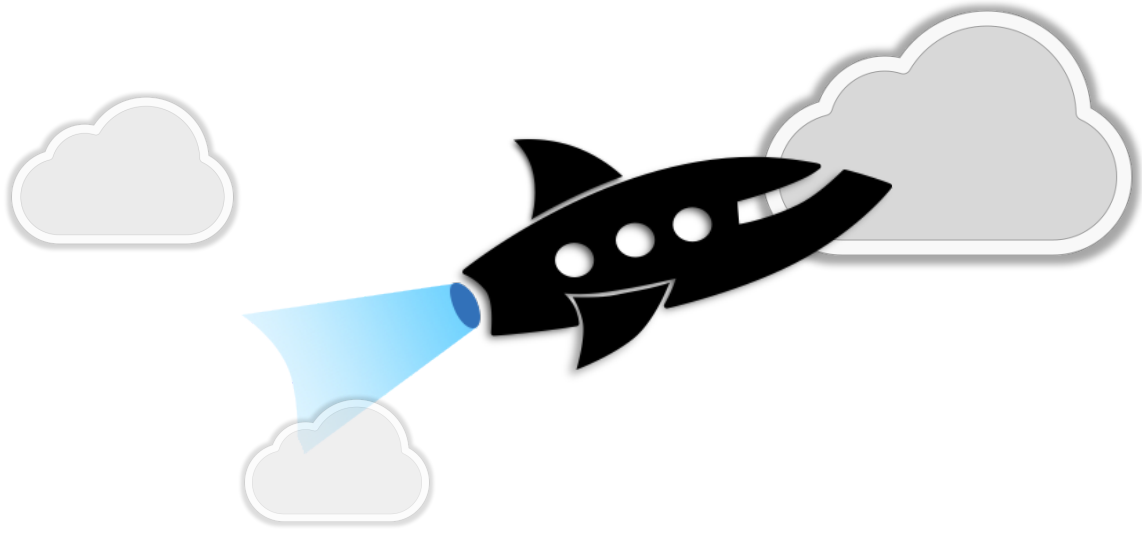

http://beyondtheclouds.github.io

anthony.simonet@inria.fr