



# Extreme Computing at **Atos**



**Case HPC Atos - CIMATEC**  
**Supercomputação aplicada a**  
**inovação, computação na borda e**  
**vídeo intelligence**

**Genaro Costa - [genaro.costa@atos.net](mailto:genaro.costa@atos.net)**

1

## About Atos

# Atos ID card

**110,000**  
experts worldwide

**Leading pure  
player in Europe<sup>1</sup>**  
Top 5<sup>2</sup> digital world leader

**€1.2bn**  
2019e operating  
margin

**€11bn**  
2019e revenue



**Worldwide IT  
Partner**  
Of the Olympic and Paralympic  
Games



Present in  
**73 countries**

**c. 5,000 patents**  
**c. €250m R&D**  
per annum

Note: Data as of December 31, 2018

1) First European player 2) Pro-forma Syntel acquisition

# Atos is organized into five divisions

---

Together, they provide the expertise to support your digital transformation

## Infrastructure and Data Management (former M/S)

We create business benefits through intelligently managed IT and digital services.

## Business & Platform Solutions (former C&SI)

We transform strategic approaches to technology, combining innovative solutions with established ones.

## Big Data & Cybersecurity

We create competitive advantage for our clients from Big Data, and ensure data is delivered safely and securely to the right parties.

## Worldline



Through Worldline, we provide unrivalled leadership in expert solutions for the fast-changing payment and digital services market.

## Unify Software & Platforms

Through Unify, we combine voice, data and video, to help our customers benefit from a unique collaboration experience.

# Big Data & Security

Atos works with organizations in the private and public sectors to generate value from their growing volumes of data. Big Data & Security unites Atos' global R&D efforts to grow its IP capabilities and portfolio. The Big Data & Security division **builds the insight & cognitive platforms of tomorrow.**

	<b>BIG DATA &amp; HPC</b> <b>Unleash the value of data</b> Act of processing, analyzing, transforming information and data The high performance computing platforms, software appliances and services
	<b>MISSION-CRITICAL SYSTEMS</b> <b>Ensure defense and intelligence</b> Critical solutions for national security, defense and aerospace, Critical systems (command and control, emergency management, homeland security, electronic systems and reconnaissance) for highly sensitive industries and government entities
	<b>CYBER SECURITY</b> <b>Make trust a business lever</b> Data and information systems protection Cyber-security products (IoT, Identity & Access Management and cryptographic products) and security services

Advanced capacity to collect, aggregate and clean data from billions of sensors in real time




End-to end security  
to protect the data,  
the process and the action value chain

Extreme Computing power to analyze and create meaning from Exabytes of information



Advanced software for automated or human decision and prescriptive insight to action



2

The need for speed ...  
computer modeling

# A bit of history...how data is transforming innovation






Old days



Current



Future



3

---

Some cases on CIMATEC




# ECG SIGNALS CLASSIFICATION

## Use of:

- Multilayer Perceptron
- Convolutional Neural Network
- Residual Networks
- Fully Convolutional Networks
- Ensemble

**Classify ECG signals from 12 sensors into 2 classes (normal and abnormal) using Brazilian dataset**



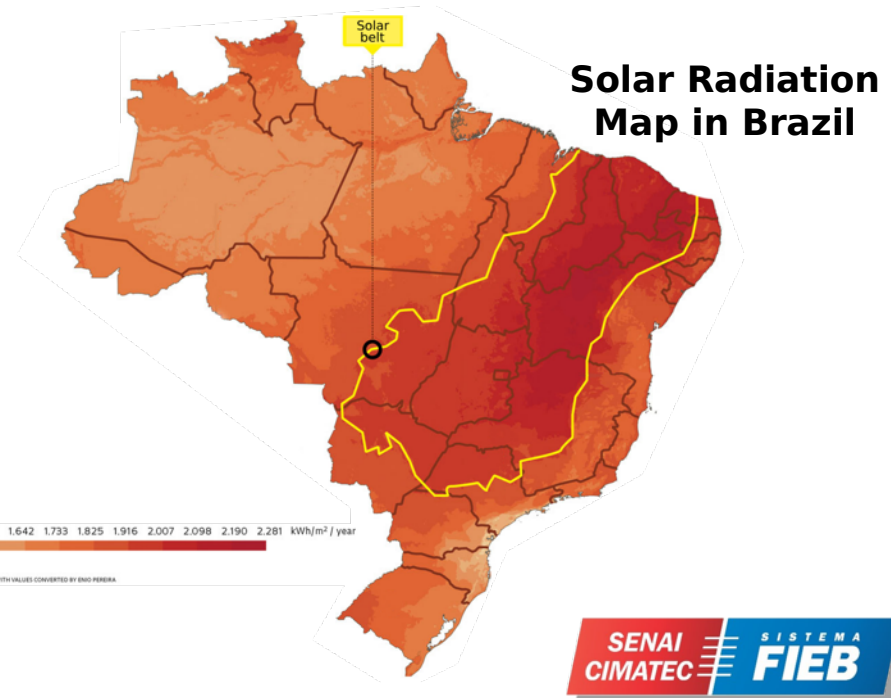
# SOLAR ENERGY PREDICTION

## Use of:



- Multilayer Perceptron;
- Boosting and traditional regression algorithms;
- Recurrent networks;
- Statistical analysis.

Analyze whether is viable or not to implement solar power plants in certain regions


Short and long range solar energy forecasting



# INITIAL VELOCITY MODEL ESTIMATION




**Input**  
Finite-differences  
seismic modeling  
data



**Target**  
Ground-truth model




**Fully Convolutional  
Network**




**Output**  
Estimated model

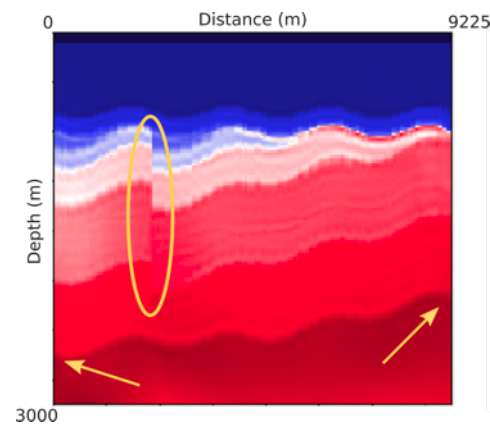
# INITIAL VELOCITY MODEL ESTIMATION




Ground-truth



FWI input



FWI output



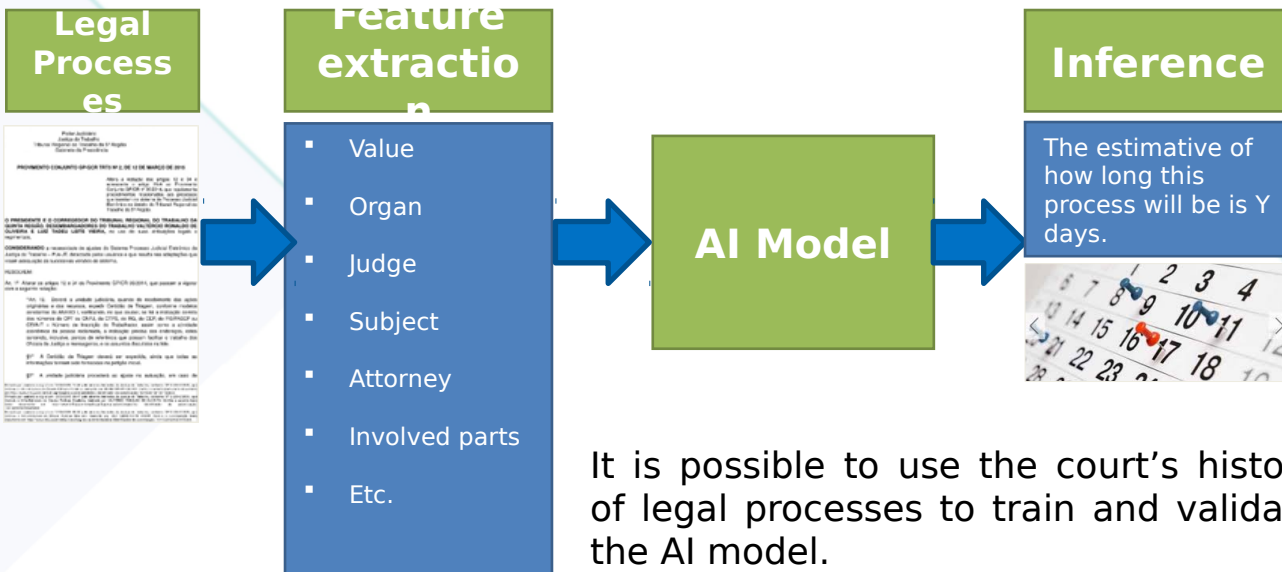
FWI's RMS Error

# DEEP LEARNING APPLIED TO LEGAL DATA



Starting date of the legal process



What will be the date when this process is finished?



# DEEP LEARNING APPLIED TO LEGAL DATA



# DEEP LEARNING APPLIED TO LEGAL DATA


Anonymous Process ID	Estimated time (days)	Real time (days)
Process 1	2.275	2
Process 2	83.525	84
Process 3	119.875	121
Process 4	49.600	50
Process 5	6.425	6
Process 6	72.800	72

# AUTOMATIC FAULT DETECTION IN ROTATING MACHINES USING DEEP LEARNING

## Use of:

- MAFAULDA dataset
- Fourier transform
- 1D CNN

Classify faults based on machines their frequency of operation





# AUTOMATIC FAULT DETECTION IN ROTATING MACHINES USING DEEP LEARNING

## General results:

- 97% when considering 30 classes
- 99% when considering 10 classes


Classes	Accuracy	Precision	Recall	F1-Score
No fault	99.98%	99.43%	100%	99.71%
Horizontal misalignment	99.94%	99.85%	99.71%	99.78%
Vertical misalignment	99.94%	99.90%	99.81%	99.86%
Unbalance	99.98%	99.92%	100%	99.96%
Sphere fault	99.96%	99.44%	99.44%	99.44%
Sphere fault and unbalance	99.98%	99.79%	100%	99.89%
Cage fault	99.98%	99.46%	100%	99.73%
Cage fault and unbalance	99.72%	98.81%	98.41%	98.61%
Outer lane fault	99.96%	99.40%	99.40%	99.40%
Outer lane fault and unbalance	99.72%	98.50%	98.50%	98.50%



# About the data volume and location

# This is just the beginning of data revolution


By 2030, data produced will reach 1 YottaByte



Today, less than 5% of all accessible data is analyzed



By 2022, data produced outside datacenter and cloud will reach 75%






By 2020, 99% of enterprise-captured video/image will be analyzed by machines



# IoT is driving the “Move to Edge”

The pendulum swings back to distributed architectures



## The Edge will eat the Cloud

*40% of large enterprises will be integrating edge computing in projects by 2021*

*-Source Gartner Inc, Sept 2017*

# 4 challenges to leverage the data revolution

## Data center platforms and appliances



- Performance
- In-memory
- Massive data sets

**1 YB**

## Trust in data



- Data quality
- Data security
- Data privacy

**5%**

## Data science AI platform



- Develop models
- Deploy from DC out to the edge
- Train & Infer models

**99%**

## Edge Computing



- Real-Time Video/Image analytics
- Edge Datalake for IoT


**75%**

# Data produced and processed (2018)

---

**Edge / factory floor / public space**

**Datacenter / Cloud**




# Data produced and processed (2022)

**Edge / factory floor / public space**

**Datacenter / Cloud**

Around 10% of enterprise-generated data is created and processed outside a traditional centralized data center or cloud. By 2022, Gartner predicts this figure will reach 75%”



<https://www.gartner.com/smarterwithgartner/what-edge-computing-means-for-infrastructure-and-operations-leaders/>

# IOT & Machine Learning enables ..






# Why do we need to bring computing to the Edge ?

Edge / factory floor / public space

Datacenter / Cloud




**Due to the rapid growth of IOT devices -  
more and more data needs to be moved from the Edge to the DC/Cloud**

# Why do we need to bring computing to the Edge ?

Edge / factory floor / public space

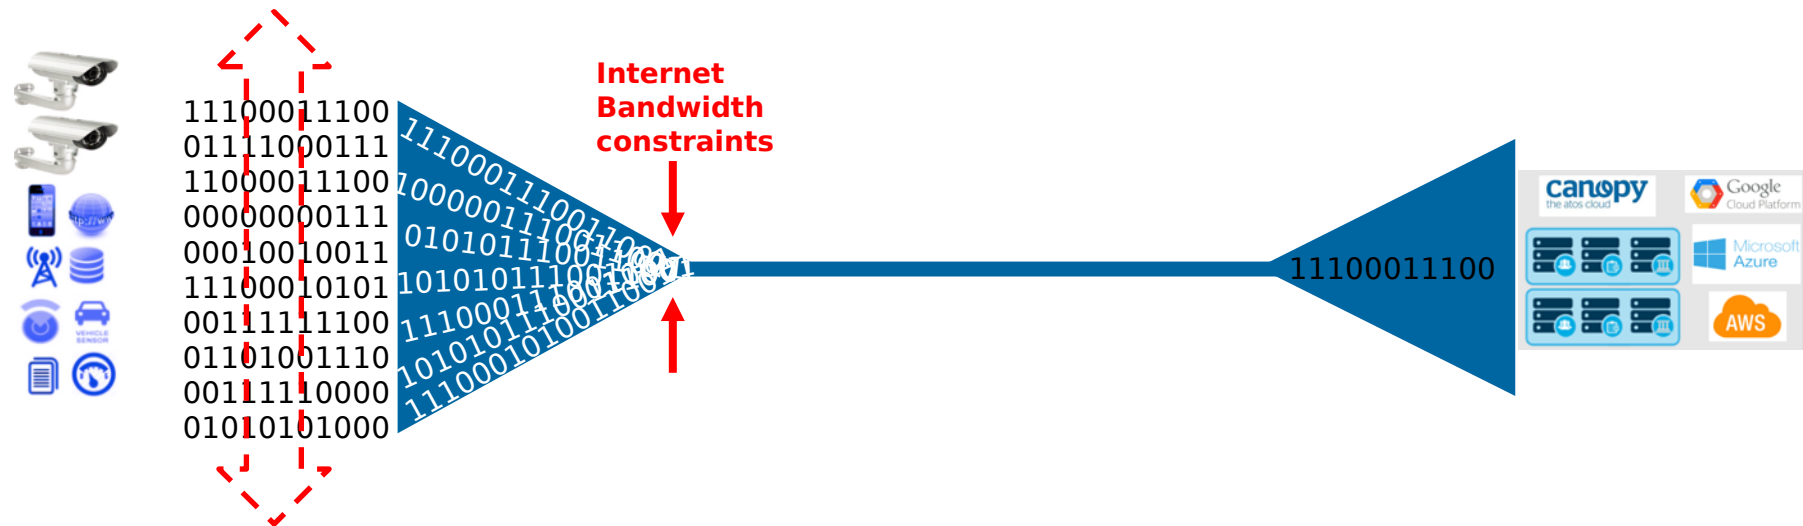
Datacenter / Cloud



# Why do we need to bring computing to the Edge ?

Edge / factory floor / public space

Datacenter / Cloud

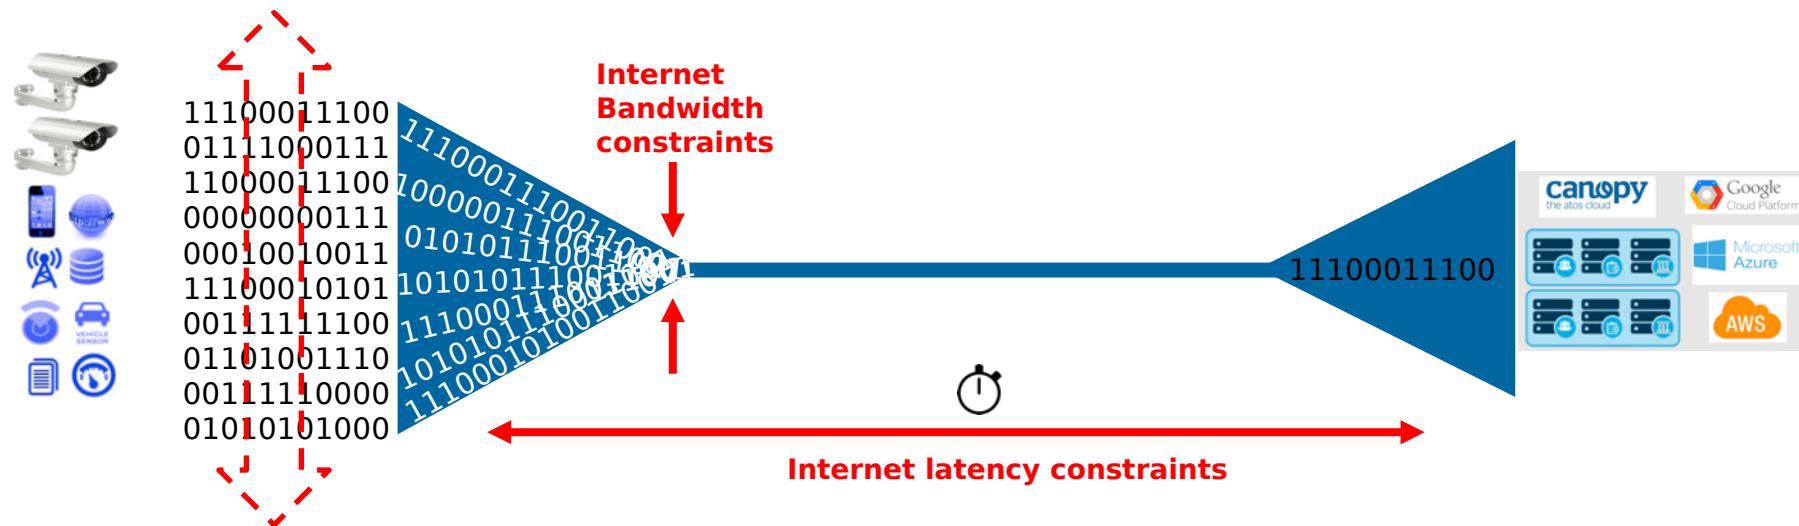


**The current infrastructure induces bandwidth and latency constraints**

# Why do we need to bring computing to the Edge ?

Edge / factory floor / public space

Datacenter / Cloud




The current infrastructure induces bandwidth and latency constraints

# Why do we need to bring computing to the Edge ?

---

**Solving the latency problem**




**Invent a solution to  
travel faster than light**




**or**

# Why do we need to bring computing to the Edge ?

**Solving the latency problem**




**Invent a solution to travel faster than light**



**or bring computing closer to the Edge**

# Why do we need to bring computing to the Edge ?



3


The Atos Edge server

---




# Atos AI-Edge Server

- ▶ **Very powerful CPU and GPU capabilities for an Edge class server**
  - designed to provide exceptional Machine Learning inference performance
- ▶ **Designed to operate outside of a Datacenter**
  - Factory/Shop floor, Airport halls, Ships, ..
  - DIN railmount option
- ▶ **Can also be mounted in a standard 19' Rack**
  - 2U form factor



# Atos AI-Edge Server

## technical specs



### ► Rugged design

- 2U form factor
- rackable & DIN rail compatible
- 5-45°C

### ► Ethernet

- 2x1Gb RJ45
- 2x 10GbE SFP+ (optional)

### ► Internal storage

- 2 SATA SSD/HDD support

### ► Security:

- TPM 2.0,
- Intel QAT , ..

### ► Memory type:

- DDR4 2133-2666
- LRDIMM

### ► Ports:

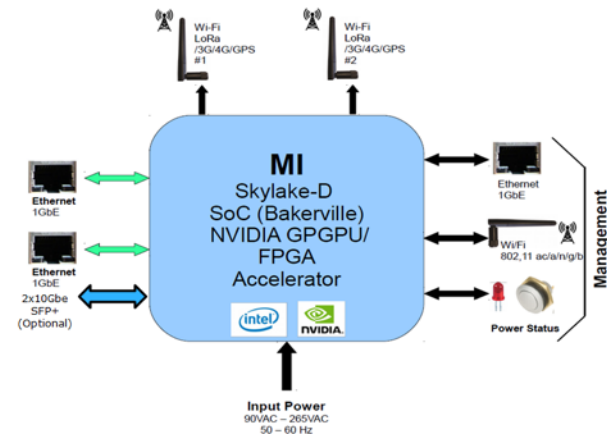
- 2x USB3

### ► OS support:


- Linux (RH), Windows

### ► Open System management

- OpenBMC
- secured remote Firmware updates



# Atos Edge Server packaging



## ► Environmental conditions

- slightly relaxed ETSI EN 300 019 class 3.2 specs, +5°C to +45°C (moving up from -5 °C to +5°C)
  - ETSI EN 300 019, Environmental conditions and environmental tests for telecommunications equipment, class 3.2: weather protected locations - Partly temperature-controlled locations.



Use cases:  
Machine learning inference

---

# Atos Edge server - Video stream analytics

exceptional ML inference performance due to accelerators

Edge / factory floor / public space

Datacenter / Cloud


Class1


Class3



**IP cameras**  
(up to 15)  
video stream



 **Codex AI Suite**  
Video intelligence  
ML model inference



**Atos**  
**Edge server**  
-up to 2 Nvidia P4 GPUs  
-or one 150W FPGA

 **Codex AI Suite**

ML model training

ML model distribution



**Atos**  
**Sequana /**  
**Escala**

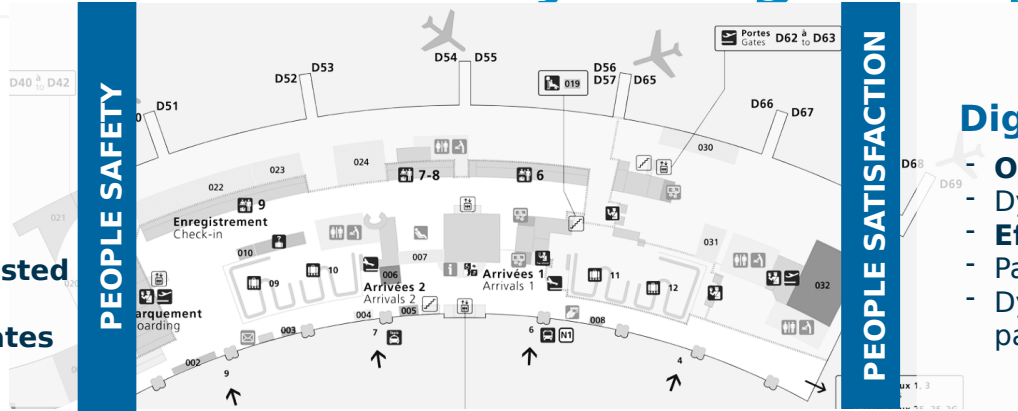
# Atos Edge server - Video stream analytics

## *AI powered Video Security & Digital Signage*

### Video Security

- Crowd **movement**
- Scenes of **violence**
- **Abandoned** objects
- Identification of **blacklisted** person
- Identification of **car plates**

PEOPLE SAFETY





PEOPLE SATISFACTION

### Digital Signage

- **Optimize** commercial spaces
- Dynamic **advertisement**
- **Efficient** signage
- Passenger **traffic** flows well
- Dynamic and **personalized** passenger information

## UNIQUE PLATFORM



# 5

## How to train the models

---





**CIMATEC AIRIS**

**500Tflop**

**1 PB -25GB/s**


**Atos**



# BullSequana X1000 cell

Compute Rack  
"extension"

# Atos




# BullSequana XH2000 / X1000

Give Access to the Best of Breed Technologies

# Atos




All In One compute rack design  
96 nodes building block  
Topology options based on customer needs  
Enhanced Direct Liquid Cooling  
InfiniBand Mellanox HDR - HDR100




~10% to ~25% cost saving compare to previous generation


Intel® Xeon®  
Processors



AMD® Rome®



**Nvidia® V100** Cavium® ThunderX2



BXI 1.2  
Fat Tree

High performance  
interconnect network

InfiniBand **HDR**  
**GHC, Fat Tree, DF+** Mellanox  
TECHNOLOGIES



# By Cluster Size

And cooling method

**20 to 1000 nodes**

## **Commodity**

X400 offer

- + **Lowest price**
- + **Good Enough**
- **Not Water Cooled**

**100 to 2500 nodes**

## **MidRange**

X550 & Bull S offers

- + **Highest Density**
- + **Cost Competitive**
- + **Ease of use**
- **Not Water Cooled**

**500 to 10000+ nodes**

## **HighEnd**

Bull Sequana X1000 offer

- + **Lowest TCO**
- + **Water Cooled**
- **Higher Cost**

6

CIMATEC  
REFERENCE CENTER ON  
ARTIFICIAL INTELLIGENCE

---

# Reference Center on Artificial Intelligence

## SENAI CIMATEC



# Trainin a

Residency

Courses

Postgraduate

Masters

Doctorate

# Project S

Executor

Moderator

Marketplace

# Servic es

Workshop

Hackathon

Operational

**We are hiring!!!**



# Thanks

---

For more information please contact:

M+ 55 71 98780-7877

genaro.costa@atos.net

Atos, the Atos logo, Atos Codex, Atos Consulting, Atos Worldgrid, Bull, Canopy, equensWorldline, Unify, Worldline and Zero Email are registered trademarks of the Atos group. June 2017. © 2017 Atos. Confidential information owned by Atos, to be used by the recipient only. This document, or any part of it, may not be reproduced, copied, circulated and/or distributed nor quoted without prior written approval from Atos.

**Bull**  
atos technologies

# CIMATEC: Verticais de Ação

## Capacitação

Residência

Cursos

Pós  
Graduação

Mestrado

Doutorado

## Projetos

Executor

Mediador

Marketplace

## Serviços

Workshop

Hackathon

Operação



# Vertical: Capacitação

## Residência

- Internatos de 4h e 8h de profissionais em IA
- Trabalho com projetos reais
- Currículos compacto incluindo da coleta a geração cognitiva
- Metodologia de desenvolvimento ágil

## Cursos

- Formação Prática em laboratório
- Verticais como

## Pós-Graduação

- Saúde/Medicina/Marketing/operação de
- Formação Latu Sensu
- Especialização ciência de dados
- Já em operação (2ª turma)

## Mestrado

- Formação Stricto Sensu
- Projeto da Indústria
- Bolsas financiadas pela Indústria

## Doutorado

- Desafios transversais
- Bolsas pagas pela indústria

# Vertical: Projetos

## Executor

- Projeto desenvolvido e gerenciado pelo CIMATEC com recursos próprios
- Equipes Multidisciplinares
- PMO
- Pool de recursos = custos otimizados

## Mediador

- Requisitos levantados pelo CIMATEC
- Gestão de projetos feito pelo CIMATEC, executado por startups/empresas.
- HUB de parceiros

## Market Place

- Entendimento e sizing da demanda buscada
- Caracterização de potenciais executores
- Análise de entregáveis

# Vertical: Serviços

## Workshop s

- Uso dos profissionais de competência sobre o domínio de negócio
- Criação de casos de negócio
- Mapa de oportunidade no uso da IA

## Hackathl on

- Times para desenvolvimento In-Loco de protótipos de prova de valor
- Uso de metodologias ágeis

## Operação

- Reuso de soluções desenvolvidas
- Acompanhamento de implantações
- Otimização continuada (treinamento/sintonização de modelos)