

<u>Best  
Practices  
Building  
Resilient  
Systems</u>



"x": 12.43

"y": 08.25

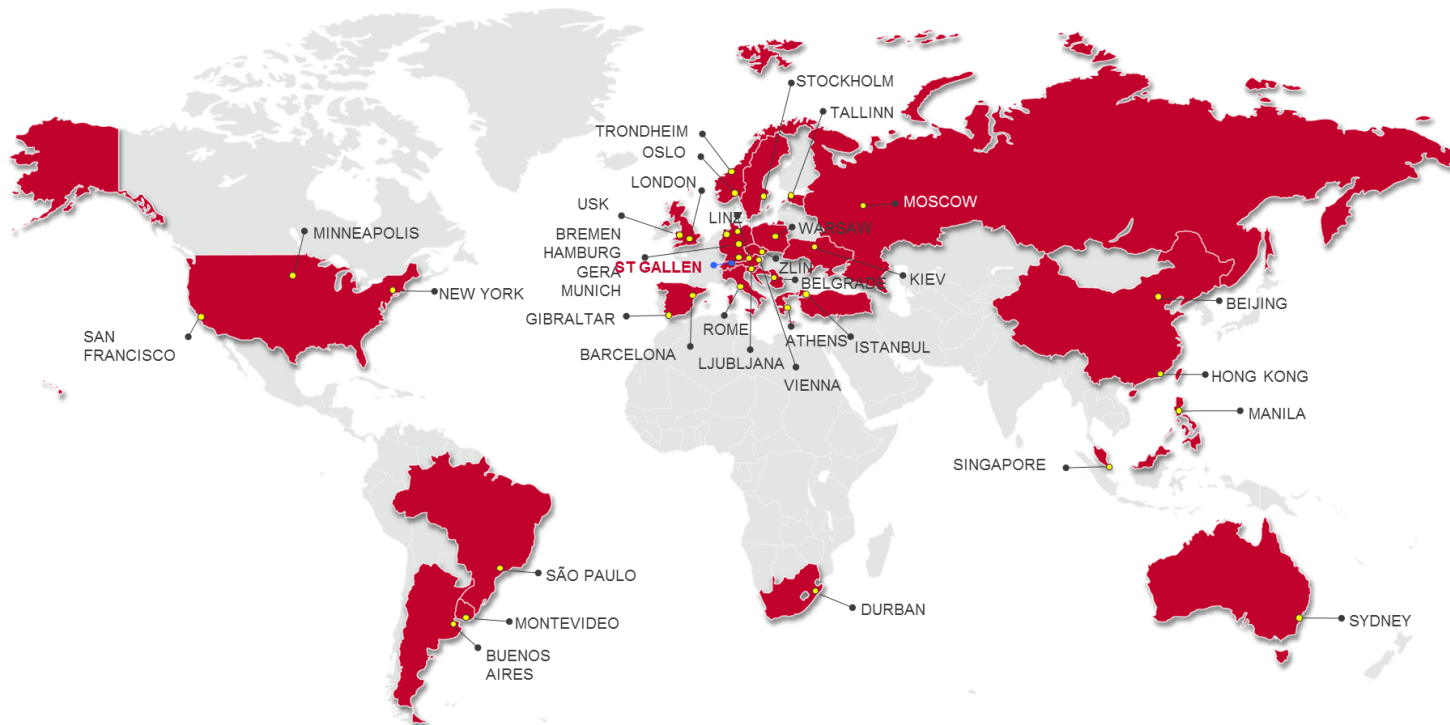
Pablo Jensen, CTO

## Who is Pablo Jensen?

- Danish – but born in Argentina where they didn't had Paul on their whitelist of names so my parents had to call me Pablo
- Computer Science degree from Copenhagen University – and MBA from Henley
- Several years in Thomson Reuters in Scandinavia, London and Switzerland
- Joined Sportradar as CTO in 2013 when the business had 500 employees with 150 in IT – now 2.000 employees and 400 in IT
- Industrial advisor for EQT
- Running, wine, car's, Brøndby IF

# Who is Sportradar?

*Operating at the intersection of sports, media and entertainment.*



- Global leader in live sports data solutions for digital sport entertainment
- 8,000+ staff and contractors globally
- 30+ global offices
- Deep coverage of more than 40 sports and 600,000 **live** events per year
- 9.000 data points updated every second
- 1 second delay from live stadium event to when data is out at our customers
- Platform handling 200,000 requests a second, serving users with up to 4gbit/s in total traffic
- 9.000 requests/second in average
- 800+ Clients and Partners

# Serving More Than 800 Global Customers



Betting





Sports Media



Integrity



Rights Holders

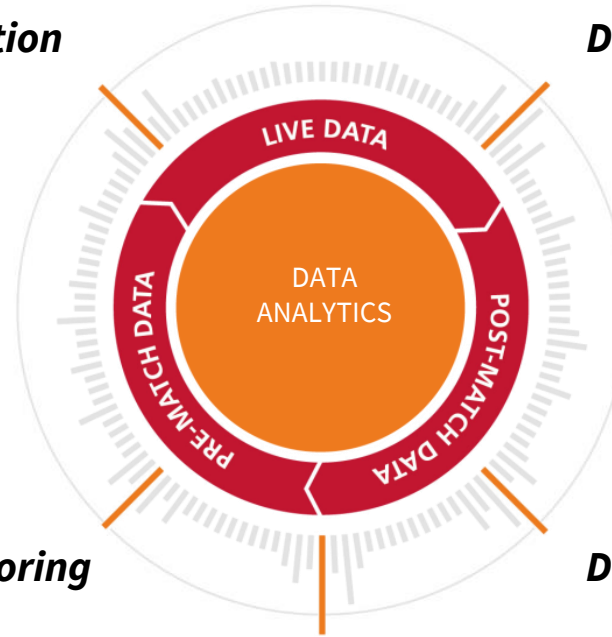


# Sportradar in a Nutshell

sp<sup>o</sup>rt<sup>r</sup>adar



**Data Collection**



**Data Processing**



**Data Monitoring**

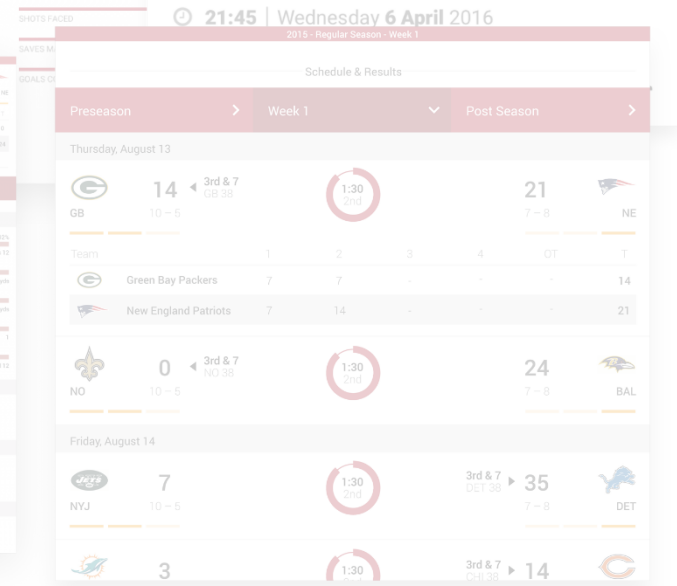

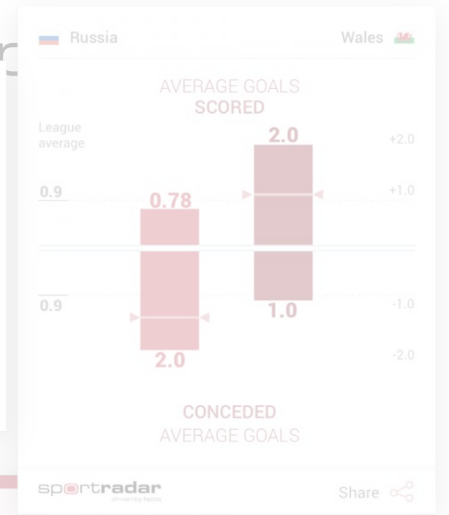
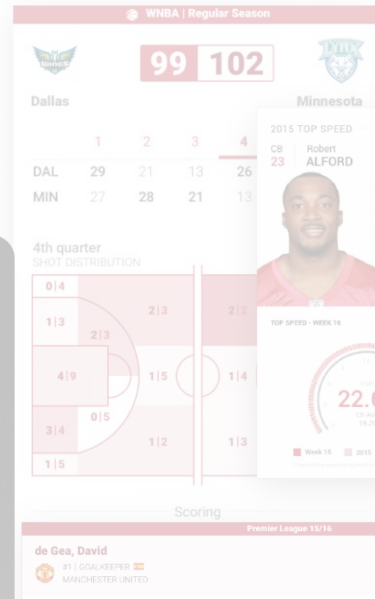
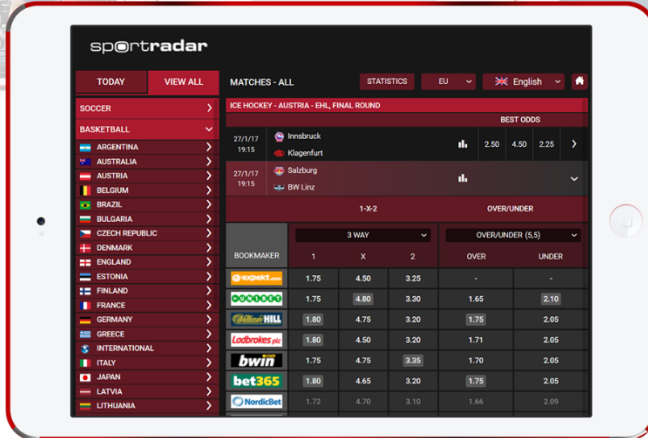

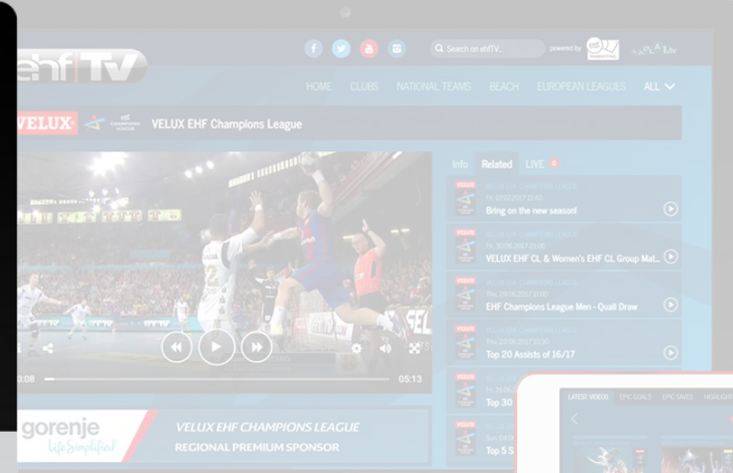
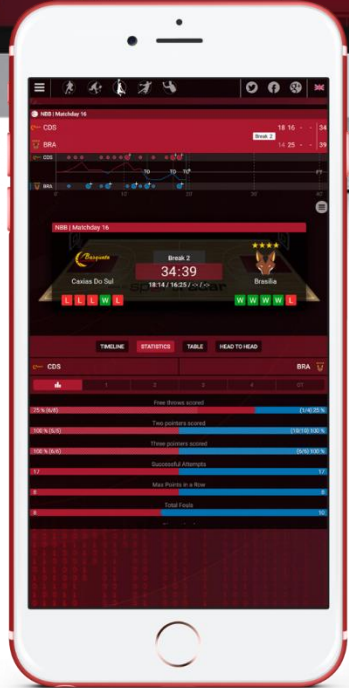

**Data Marketing**



**Digital Sports Solutions**



# Sports Media: Live Score



# Sports Media: AV & OTT

The image displays a variety of sports media interfaces across different devices:

- Desktop Monitor:** Shows a live match tracker for a football game with a 3:2 score and a video player for a VELUX EHF Champions League match. The interface includes navigation menus, social media links, and a list of related content.
- Tablet:** Displays a grid of video thumbnails for various sports events, including highlights and full matches.
- Smartphone (Top Left):** Shows a live match tracker for a football game with a 3:2 score and a video player.
- Smartphone (Bottom Left):** Displays a detailed statistics dashboard for a football match, including possession, shots, and player ratings.
- Smartphone (Bottom Center):** Shows a detailed statistics dashboard for an ice hockey match, including goals, assists, and player ratings.
- Smartphone (Bottom Right):** Displays a detailed statistics dashboard for a football match, including possession, shots, and player ratings.
- Tablet (Top Right):** Shows a detailed statistics dashboard for a basketball game, including player stats, team performance, and game flow.
- Tablet (Middle Right):** Displays a detailed statistics dashboard for a football match, including possession, shots, and player ratings.
- Tablet (Bottom Right):** Shows a detailed statistics dashboard for a football match, including possession, shots, and player ratings.

# Sports Media: Widgets & Cards

The image displays a variety of sports media widgets and cards across different devices:

- Desktop Monitor (Top Left):** A 'Live Match Tracker' for a soccer match, showing a 3:2 score and a stadium view.
- Desktop Monitor (Top Center):** An 'enTV' video player for a VELUX EHF Champions League match, featuring a video player, navigation controls, and a sidebar with related content.
- Desktop Monitor (Top Right):** A WNBA Regular Season scoreboard for Dallas (99) vs Minnesota (102), including a shot distribution chart and player statistics for Robert Alford.
- Desktop Monitor (Bottom Right):** A Premier League 15/16 player profile for David de Gea, showing 763 minutes played and a 0.38 clean sheet ratio.
- Tablet (Middle Right):** A 'Schedule & Results' widget for the Chicago Bulls vs Los Angeles Lakers game, showing scores and game details.
- Tablet (Bottom Center):** A widget for the 'ICE HOCKEY - AUSTRIA - EHF FINAL ROUND' with a table of matches and betting odds.
- Smartphone (Bottom Left):** A mobile version of the live match tracker, showing a 34:39 game time and various match statistics.
- Smartphone (Bottom Center):** A mobile version of the video player, showing a smaller view of the EHF Champions League match.
- Smartphone (Bottom Right):** A mobile version of the schedule and results widget, showing game details for Thursday, August 13 and Friday, August 14.



# Betting: Life Cycle of Odds

The image illustrates the 'Life Cycle of Odds' in betting through several interconnected components:

- Main Monitor (Left):** Displays a detailed betting odds table for a horse race. A blue circular callout highlights a specific section of the table with the following values:

1.750	1.570
2.090	2.240
105	105
1.830	1.740
1.772	1.772
1.970	1.960
2.059	2.059
104	104
2.090	1.980
2.090	2.090
1.740	1.730
- Smartphone (Bottom Left):** Shows a mobile betting app interface with a list of events and odds.
- Tablet (Middle Right):** Displays a live sports broadcast of a horse race with a betting overlay.
- Large Tablet (Bottom Right):** Shows a complex betting management system with various settings and data columns, including 'Subversion', 'Event Creation & Odds Management', 'Trading Tools', 'Resolving', 'Live', 'Task', 'Configuration', 'Product View', and 'Support'. It includes a legend for match availability and a 'Manual' section for live betting controls.

# Betting: Live Odds



# Betting: Virtual Games

**FAST-PACED REAL MONEY BETTING**

**EXPAND YOUR VIRTUAL SPORTS BETTING PORTFOLIO**

**AND BOOST YOUR BUSINESS**

Meeting	Time	Win	Place	Show
Meeting 4841 - 1212	00:00	Win	Place	Show
Meeting 4842 - 412	03:30	Win	Place	Show
Meeting 4842 - 312	05:50	Win	Place	Show
Meeting 4842 - 312	08:01	Win	Place	Show
Meeting 4842 - 412	09:52	Win	Place	Show
Meeting 4842 - 112	03:30	Win	Place	Show
Meeting 4842 - 212	05:50	Win	Place	Show
Meeting 4842 - 312	08:01	Win	Place	Show
Meeting 4842 - 412	09:52	Win	Place	Show

Rank	Item	Price	Value
1	Seamus Bramble	4-1	3-6
2	Zoe Pepper	3-6	2-6
3	Hendrix Bob	1-3	1-4
4	Dash Cosmo	4-2	3-4
5	Crown Artemis	5-4	5-2
6	Sugar Sassy	3-4	5-1

Rank	Item	Price	Value	
1	Chief Mayhem	6.25	2.80	1.65
2	Climb to top	12.00	6.75	3.15
3	Prince Blanket	2.25	1.25	1.08
4	Peritas Marlowe	3.40	1.65	1.16
5	Fabulous Hans	7.00	3.05	1.70
6	Mr. Bullet	25.00	8.50	4.85

# Betting: Integrity

- FAST-PACED REAL MONEY BETTING
- EXPAND YOUR VIRTUAL SPORTS BETTING PORTFOLIO
- AND BOOST YOUR BUSINESS

The interface shows a 'Fraud Detection System' with a sidebar of sports categories like Football, Basketball, and Cricket. The main area contains a table with columns for 'Home Team', 'Away Team', 'Matchdate', 'W/L', 'Score', and 'Margin'. Below the table is a line graph showing performance trends over time.

The interface displays 'Meeting 4841 - 1212' for 'Belmore Park Stayers'. It lists race results with columns for 'Rank', 'Horse Name', and 'Time'. A photo of a horse race is shown in the background.

Rank	Horse Name	Time
1	Deamus Bramble	4:16.3-8
2	Zoe Pepper	3:62-8-4
3	Hendrix Bob	1:31-1-4-2
4	Dash Cosmo	4:23-4-1
5	Crown Athena	5:45-2-1
6	Sugar Sassy	3:45-1-5

A large, multi-column data table interface, likely for betting odds or market data. It features a grid of columns with various numerical values and text entries, possibly representing different betting markets or odds for various events.

A collection of device screens showing betting-related data. It includes a desktop monitor with a data table, a tablet displaying a list of items, and a smartphone showing a betting slip or odds screen.



A dashboard for 'betradar' showing various sports data. It includes a live match score of '0:2' and a list of betting markets with odds. The interface is clean and modern, with a focus on real-time data.

# Data Feeds & Development Services

## SPORTS API

.xml and .json sports data feeds for your application




eSports Service



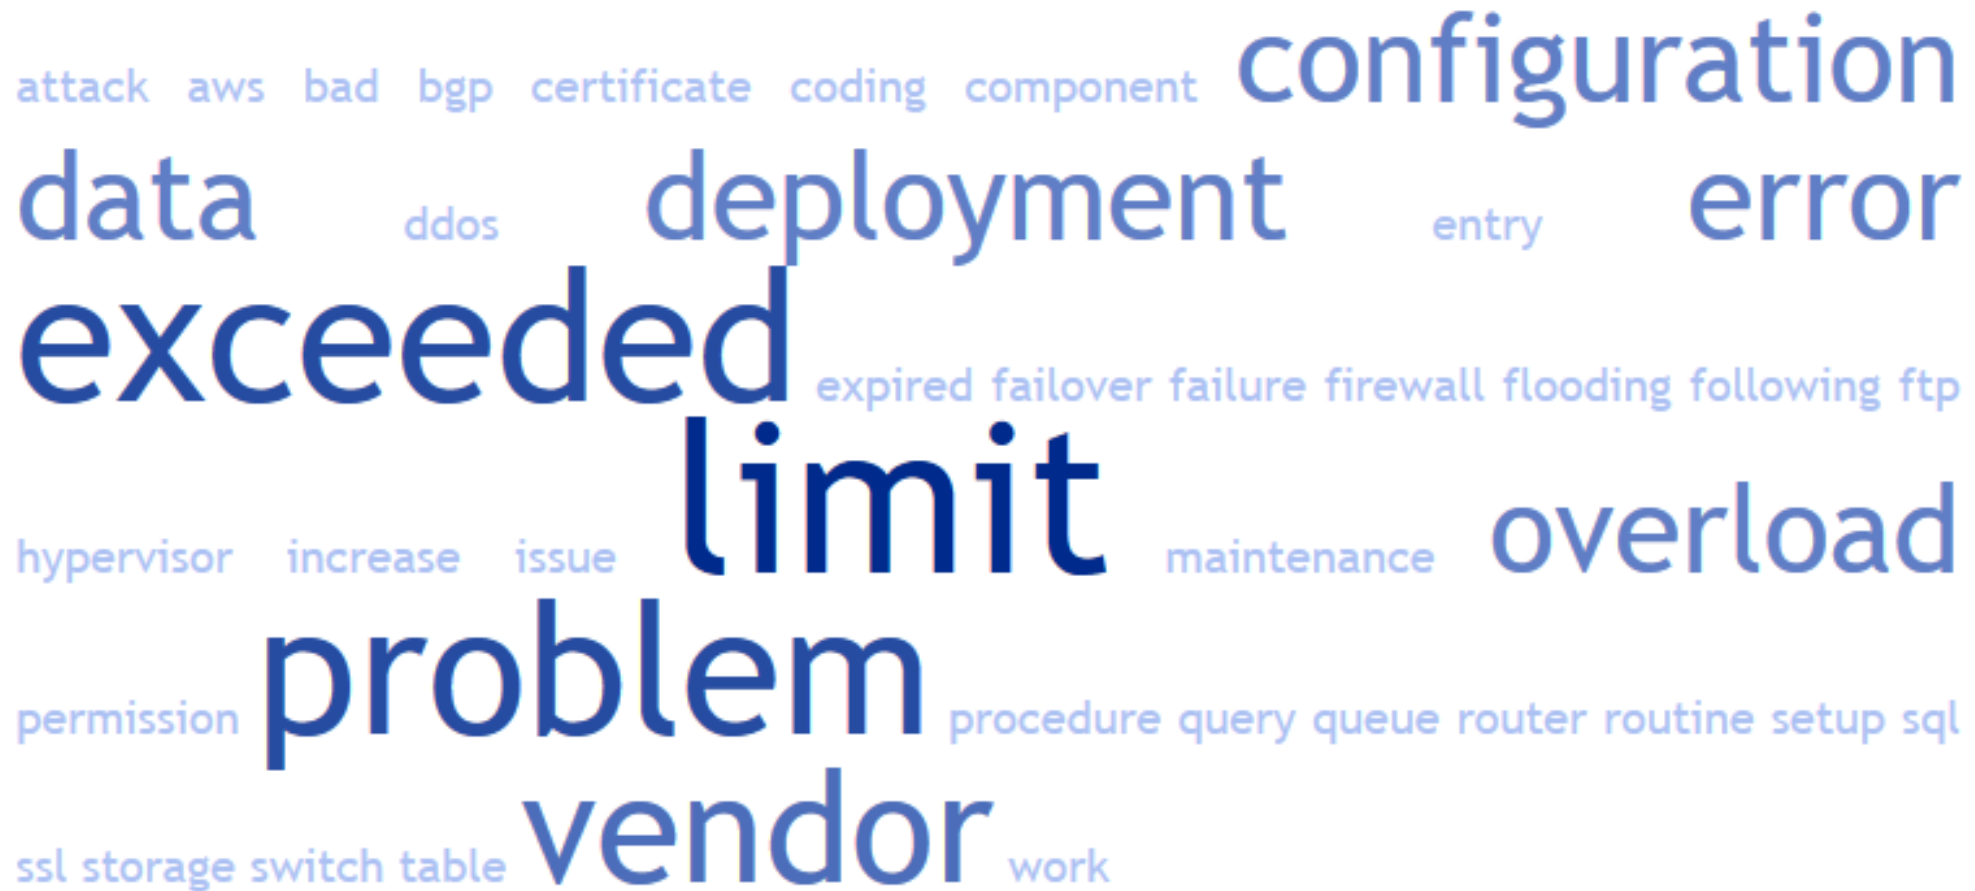
Live Odds



MDP - Mobile Development Platform



# What can go possibly wrong??



# What can go possibly wrong??

## Top incident reasons

1. 3<sup>rd</sup> Party Provider issue
2. Limit exceed (table, storage, traffic)
3. Coding error
4. Not following agreed procedures

Prepare for that there *always* will be something wrong

Physical Footprint

Technology

Process

# IT Organisation Tech Stack

## 400+ employees in 10+ IT locations:

- 40+ Dedicated teams
- 300+ Developers
- 35 Tech Leads
- 40+ System Administrators
- 40+ Project Managers
- 30+ QA
- 20+ Mobile Developers



## Web:

HTML5/CSS3, React, Javascript, API driven, Nginx, NodeJS, Varnish, Tomcat, Jetty

## Mobile:

IOS, Android

## Backend:

Java, PHP, Scala, JRuby, Go, C++, Memcache, Redis, MySql, Cassandra, MongoDB

## Sys Admin

Ganeti, OpenStack, Zabbix, Puppet, Mcollective, Debian Linux, AWS, Ceph, Kubernetes

## Source code system

GIT (GitLab)

## Open Source scanning

WhiteSource

## Build management:

Jenkins, GitLab CI

## BI & Analytics:

S3, ORC, NiFi, RedShift, Athena, Spark, Qlik

## Communication Tools

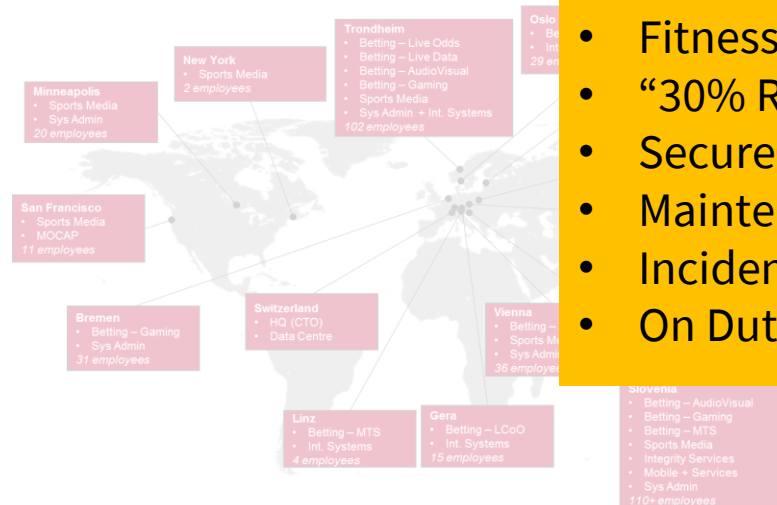
Slack, Outlook, own build tools for Incident and Maintenance Management but looking at migrating to 3rd party services (StatusPage.io)



# IT Organisation Tech Stack

## 400+ employees in 10+ IT locations:

- 40+ Dedicated teams
- 300+ Developers
- 35 Tech Leads
- 40+ System Administrators
- 40+ Project Managers
- 30+ QA
- 20+ Mobile Developers



## Web:

HTML5/CSS3, React, Javascript, API driven, Nginx, NodeJS, Varnish, Tomcat, Jetty

## Mobile:

IOS, Android

## Best practices for building resiliency

Strict defined tech stack – new technologies are architecture driven, not developer driven

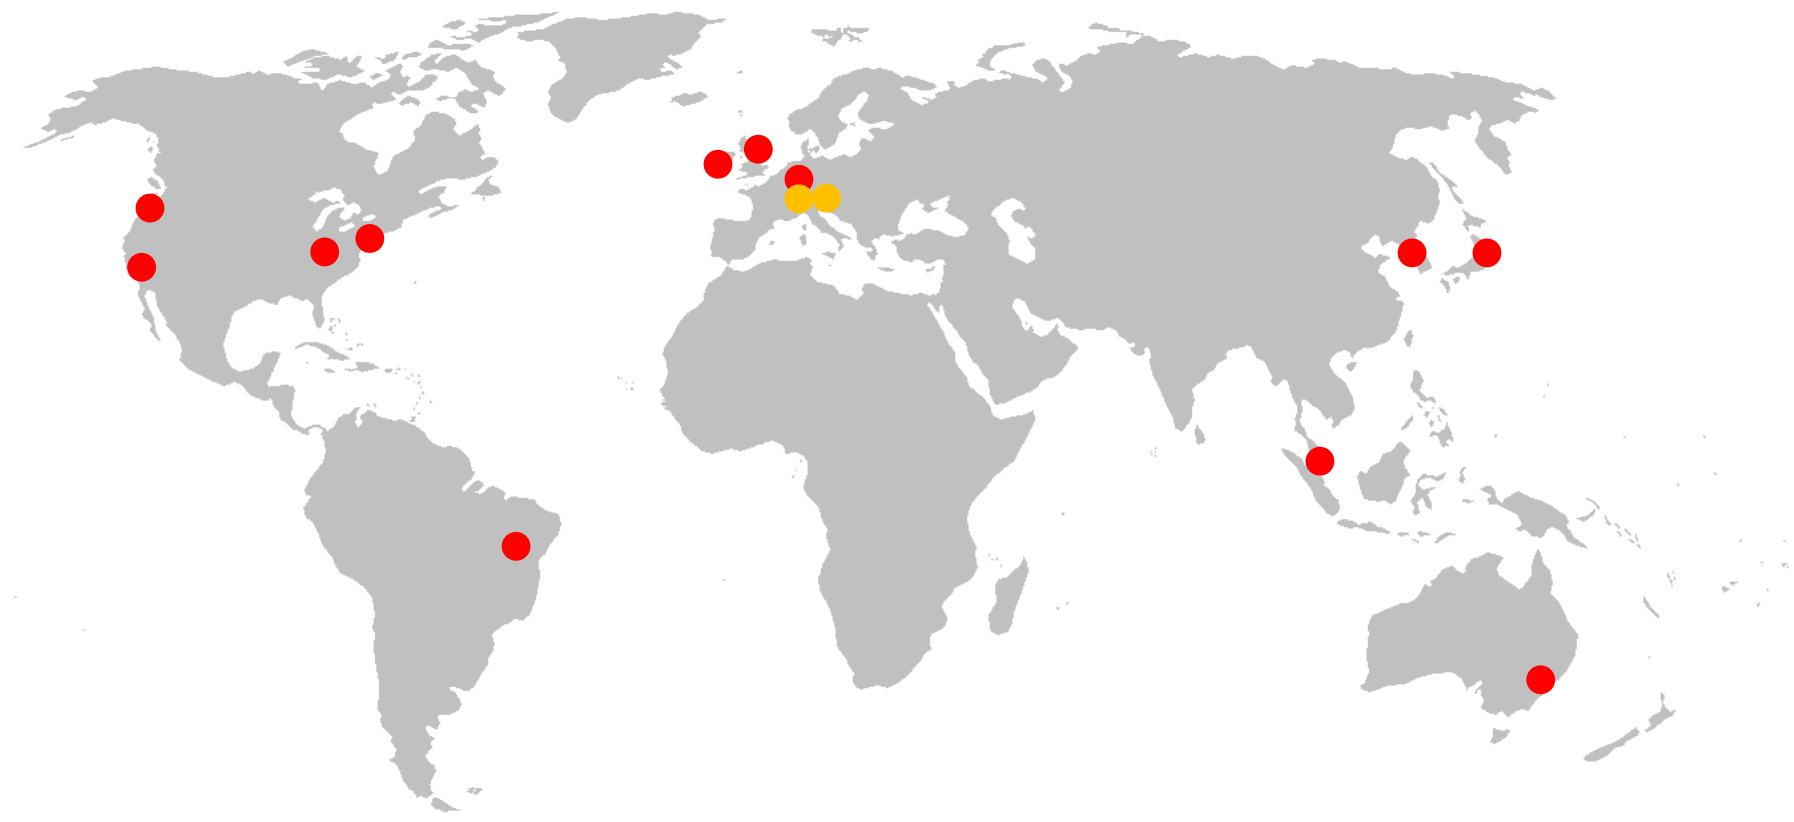
Key technical IT gate points to be followed

- Fitness for Development
- Fitness for Launch
- “30% Rule”
- Secure Development Guidelines
- Maintenance Procedure
- Incident Procedure
- On Duty Procedure

Redis, MySQL, Cassandra, MongoDB

Debian Linux, Amazon Web Services,

# Sportradar Hosting Locations



- Own regional based data center locations in Europe
- AWS/Amazon hosting locations used by Sportradar

# Sportradar Hosting Locations


**Physical Footprint**

**Best practices for building resiliency**

Identical physical regional located core data centers running live-live treated as single redundant data center.

Multiple options for client access:

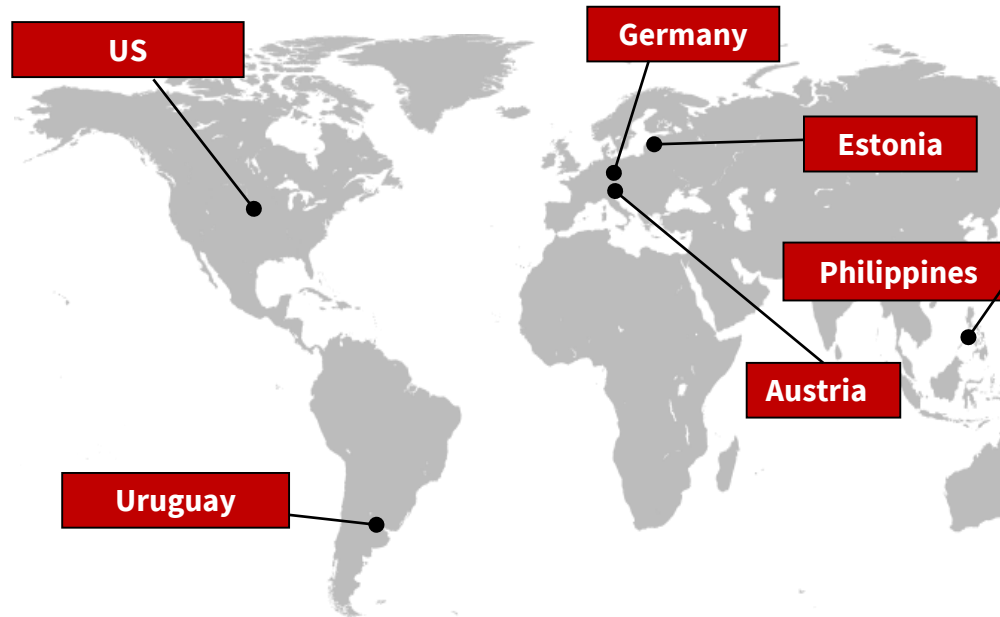


- Strategic located POP's
- Direct connect
- Open Internet



# Sportradar's Global Data Production

Sportradar Production with more than 900 employees globally

*Operations setup is physical redundant so we can shift operations between locations*



## Key facts




- Worldwide accepted data quality unmatched in combination of speed and accuracy
- Redundant production setup
- Key positions manned with branch expertise from all business segments
- State of the art data entry tools, developed in-house, enhanced based on needs of operations
- Operations approved and well-rehearsed, permanently reviewed and improved/adjusted
- >900 operators across 7 locations
- >6,000 scouts globally

# Sportradar's Global Data Production

*Operations setup is physical redundant so we can shift operations between locations*

Sportradar Production with more than 900 employees globally

## Physical Footprint



**US** **Best practices for building resiliency**

Identical production locations

Tasks can move from one location to another



Wide accepted data quality unmatched in  
 combination of speed and accuracy

Redundant production setup

Locations manned with branch expertise  
 across all business segments

- State of the art data entry tools, developed in-house, enhanced based on needs of operations
- Operations approved and well-rehearsed, permanently reviewed and improved/adjusted
- >900 operators across 7 locations
- >6,000 scouts globally

## Providers

All service elements; eg. ISP, CDN, DDOS Protection, cloud hosting, physical hosting, DNS, physical production locations, POPs, fixed line connections are understood and categorized with full risk understanding and acceptance.

# Providers

## Physical Footprint

All service elements  
 hosting, physical  
 POPs, fixed line  
 with full risk u

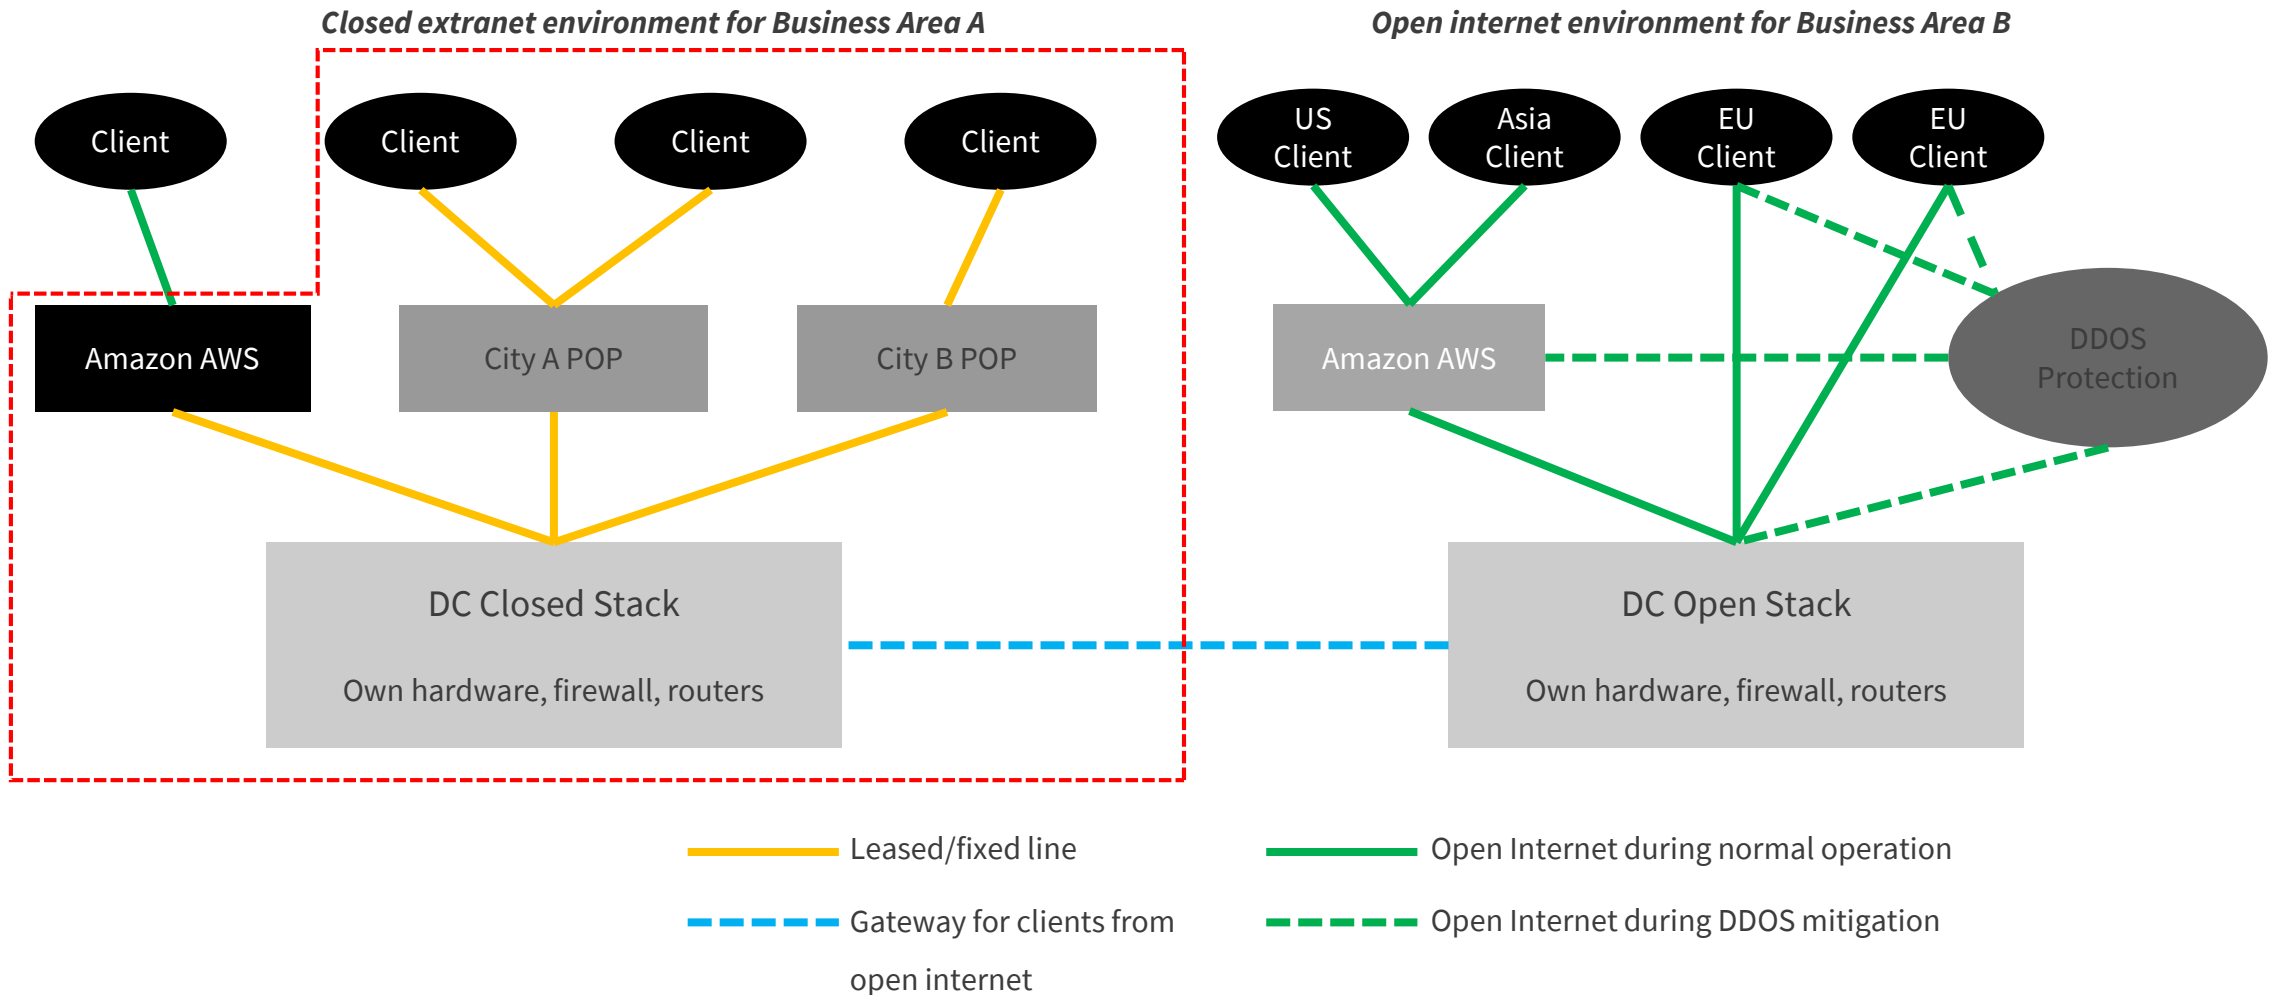
### Best practices for building resiliency

Understand and accept:

- Service elements that are 'multi-vendor'
- Service elements that are 'multi-regional'
- Service elements that are 'single' served

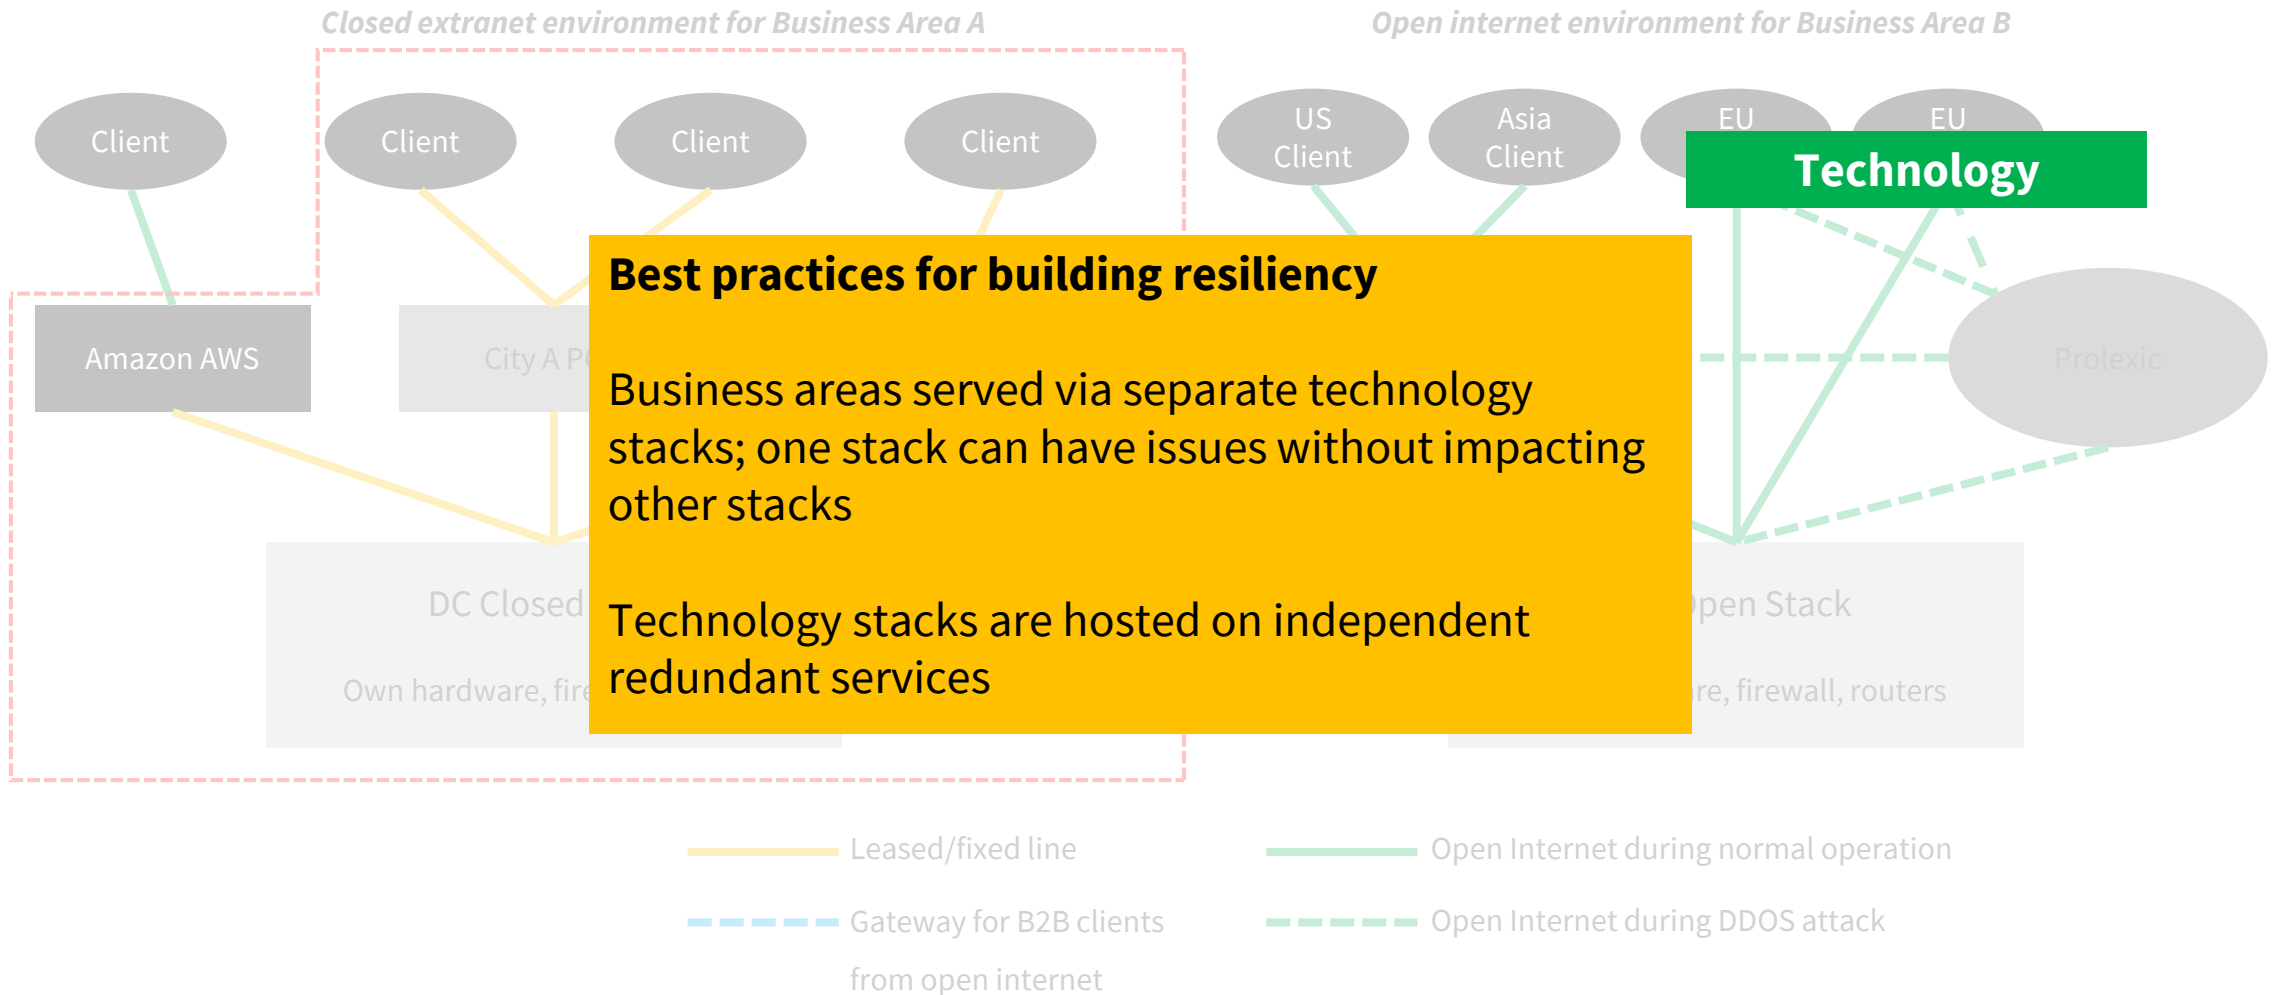
on, cloud  
 on locations,  
 categorized

# Separate technology stacks





# Separate technology stacks



# Architecture Deployment Model

## One of our Backend Core Systems

- Running on 3 dedicated physical servers in 3 different physical locations
- Composed of many sub-systems - each running as an independent cluster
- Java services either stateless or stateful while keeping data in a distributed mem-grid
- Clustered active-active setup of RabbitMQ, Zookeeper, HAProxy, Mongo replica sets, Cassandra
- Master-slave active-passive setup of MySQL, MySQL Fabric and Redis instances
- Mongo point-in-time incremental backup, MySQL/Redis/ZK daily backups
- Recovery mechanisms (e.g. a subsystem is able to recover its state based on reference data)
- Async service design (message passing, streaming)
- Circuit-breakers, request throttling, fail-fast approach (Hystrix)
- Decoupling of operational and archive/warehouse databases
- Decoupling and different types of disk volumes, reduce I/O contention (e.g. Mongo, MySQL, Backup, VMs)
- Lots of attention to low-latency implementation and design

**3 availability zones**

**Separate cluster per sub system**

**Active/Active**

**Active/Passive**

**Recovery**

**Async Design**

**Decoupling**

# Architecture Deployment Model

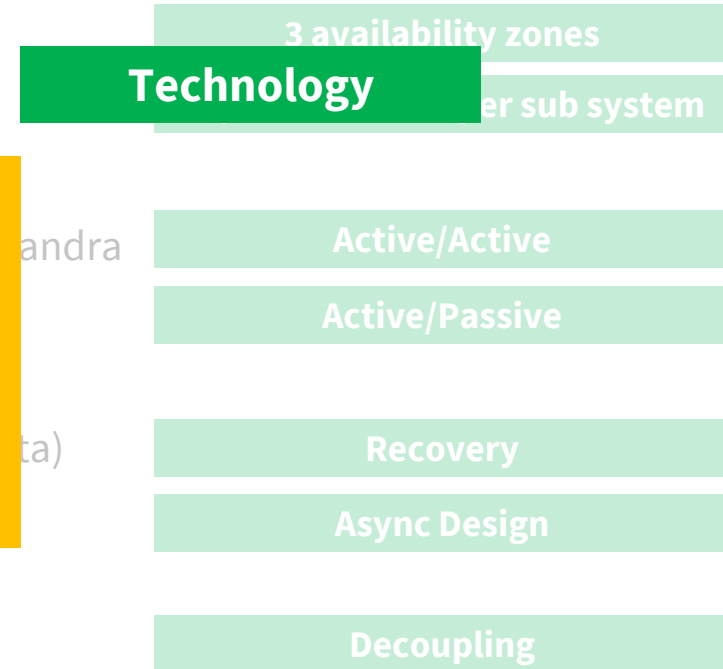
## One of our backend Core Systems

- Running on 3 dedicated physical servers in 3 different physical locations
- Composed of many sub-systems - each running as an independent cluster
- Java services either stateless or stateful
- Clustered active-active setups
- Master-slave active-passive setups
- Mongo point-in-time incremental backups
- Recovery mechanisms (e.g. Zookeeper, Kafka)
- Async service design (message queues)
- Circuit-breakers, request throttling, fail-fast approach (Hystrix)
- Decoupling of operational and archive/warehouse databases
- Decoupling and different types of disk volumes, reduce I/O contention (e.g. Mongo, MySQL, Backup, VMs)
- Lots of attention to low-latency implementation and design

### Best practices for building resiliency

Clear technical guidelines for development teams;  
no need to invent the wheel all the time

*Test, test, test*



# Architecture Deployment Model

## Frontend Systems

- Applications assume little about the infrastructure it runs on
- Can be deployed to cloud or on premise
- Servers are provisioned the same way regardless if they are on premise or in the cloud
- Conservative about using cloud or on premise services that lock us to that infrastructure (especially cloud services)
- Redundant direct connect links between on premise and cloud infrastructure
- Route53 load balance between DCs. Very handy when ISPs fail, incoming traffic just flows through other DCs that then use the direct connect backbone to reach correct destinations
- On-premise usually takes the bulk of the traffic due to traffic costs.

**Frontend / Backend Separation**

**Deployed everywhere**

**Avoid infrastructure lock in**

**Direct connections**

**Route53 to LB between DC's**

**On premise to reduce AWS costs**

# Architecture Deployment Model

## Frontend Systems

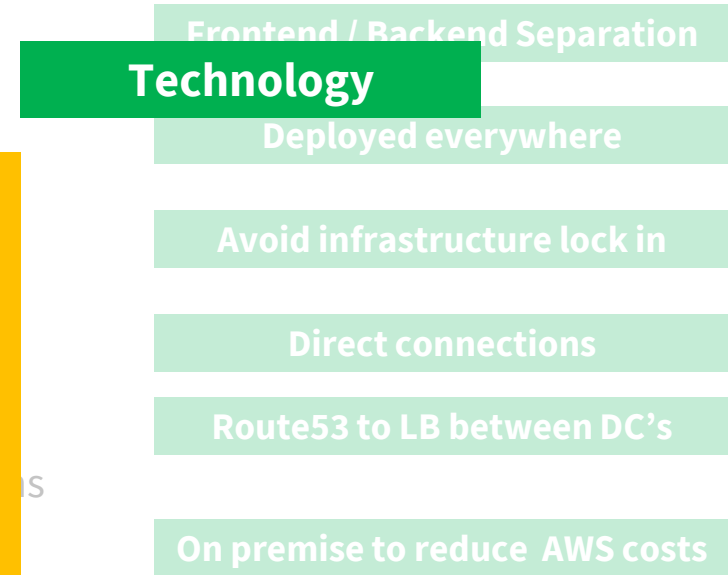
- Applications assume little about the infrastructure it runs on
- Can be deployed to cloud or on premise
- Servers are provisioned the same way
- Conservative about using cloud services (especially cloud services)
- Redundant direct connections between DC's
- Route53 load balance between DC's through other DC's that the application is running on
- On-premise usually takes the same amount of time to provision as cloud

### Best practices for building resiliency

Frontend / backend infrastructure separation to ensure no vendor lock in

Frontend technology can be deployed everywhere

Use direct connections where possible



# Client Support & Technical Support

## Setup of Client Support & Technical Support

- Global 24x7x365 support service via Chat, Helpdesk & Phone
- ISO 9001 certified
- Escalate to relevant technical On Duty Team
- All teams with a service in production are required to have a 24x7x365 On-Duty Team
- Only best engineers part of such a team
- Own build tools for on call and incident management – looking at PagerDuty

„Your reps were very professional and speedy in replying to my emails which is something I appreciate.“ – **Client A**

„Great support - very fast and accurate answers.“

**Client B**

„Quick response and solution. Re-sending the feed was everything we needed and it fixed the problems on our side.“

Thanks guys.“ – **Client C**

54%

Of all customer tickets have been **finally** solved in **less than 60 minutes**

< 0.5%

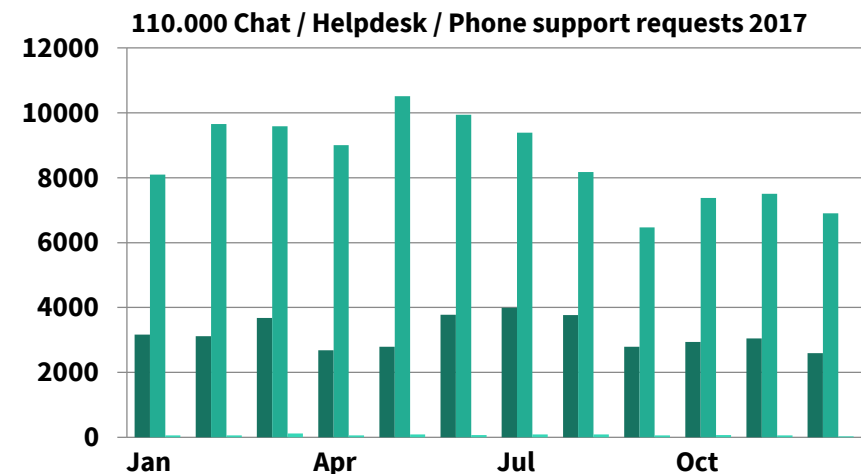
Of all incoming requests escalates to technical support

98%

Is the average handling rate of all incoming chats

97%

Incoming chats have been accepted by an operator in **less than 18 seconds**



# Client Support & Technical Support

## Setup of Client Support & Technical Support

- Global 24x7x365 support service
- ISO 9001 certified
- Escalate to relevant technical team
- All teams with a service in place
- 24x7x365 On-Duty Team
- Only best engineers part of

„Your reps were very professional and speedy in replying to my emails which is something I appreciate.“ – **Client A**

„Great support - very fast and accurate answers“

### Best practices for building resiliency

Support is Your friend – 110.000+ requests/year

Less than 0.5% escalates to a technical team

Incident reviews

Be open and transparent – also during incidents

Focus on monitoring – can always be improved – remember down to low level if you do on premise

Tight control on communication channels and on call and incident management tools

*Developers Eat Own Dog Food*

54%

### Process

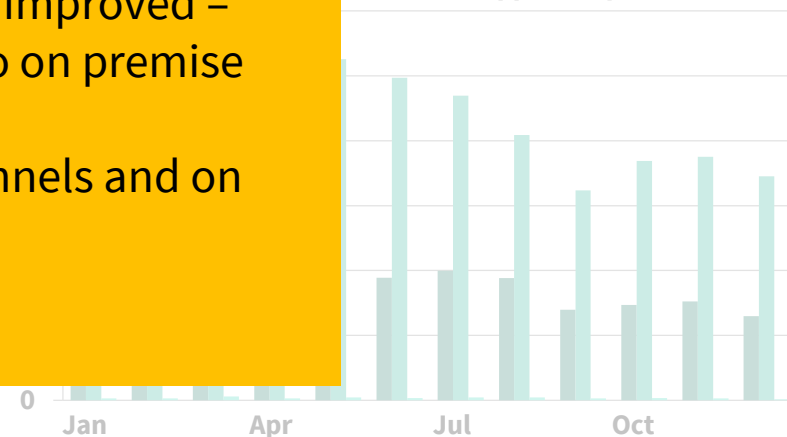
< 0.5%

Of all incoming requests escalates to technical support

97%

Incoming chats have been accepted by an operator in less than 18 seconds

esk / Phone support requests 2017



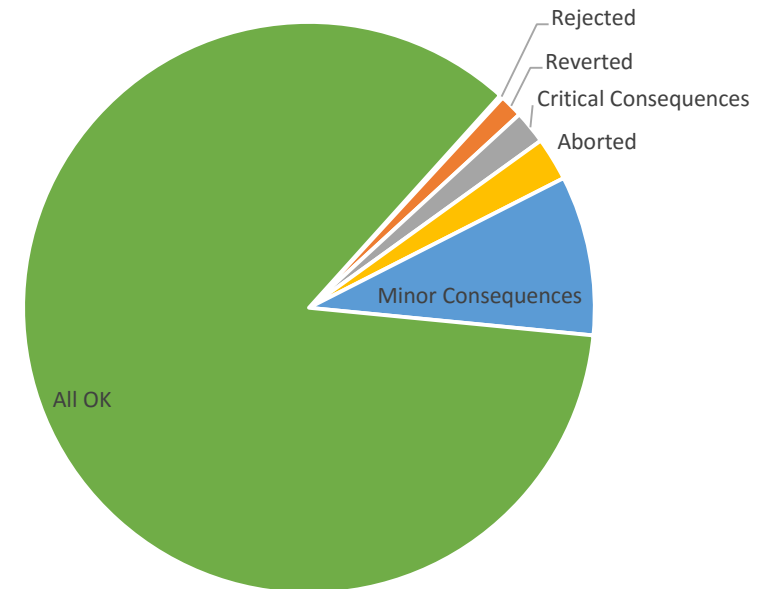

# Maintenance Procedure

## Maintenance procedure with clear rules

- Affected clients to be notified at least 2 days in advance
- Always scheduled in cooperation with Operations
- Friday/Saturday/Sunday no maintenance – if it happens then it's low risk and/or with business approval

## Good results

- Less than 4% aborted
- Close to “rolling updates”





# Maintenance Procedure

## Maintenance procedure with

- Affected clients to be notified
- Always scheduled in cooperation with
- Friday/Saturday/Sunday no maintenance then it's low risk and/or with

## Good results

- Less than 4% aborted
- Close to "rolling updates"

## Best practices for building resiliency



Clear rules – common for all teams

Need for exceptions build into the process

Respect your peak days

Ensure strong tools for process management

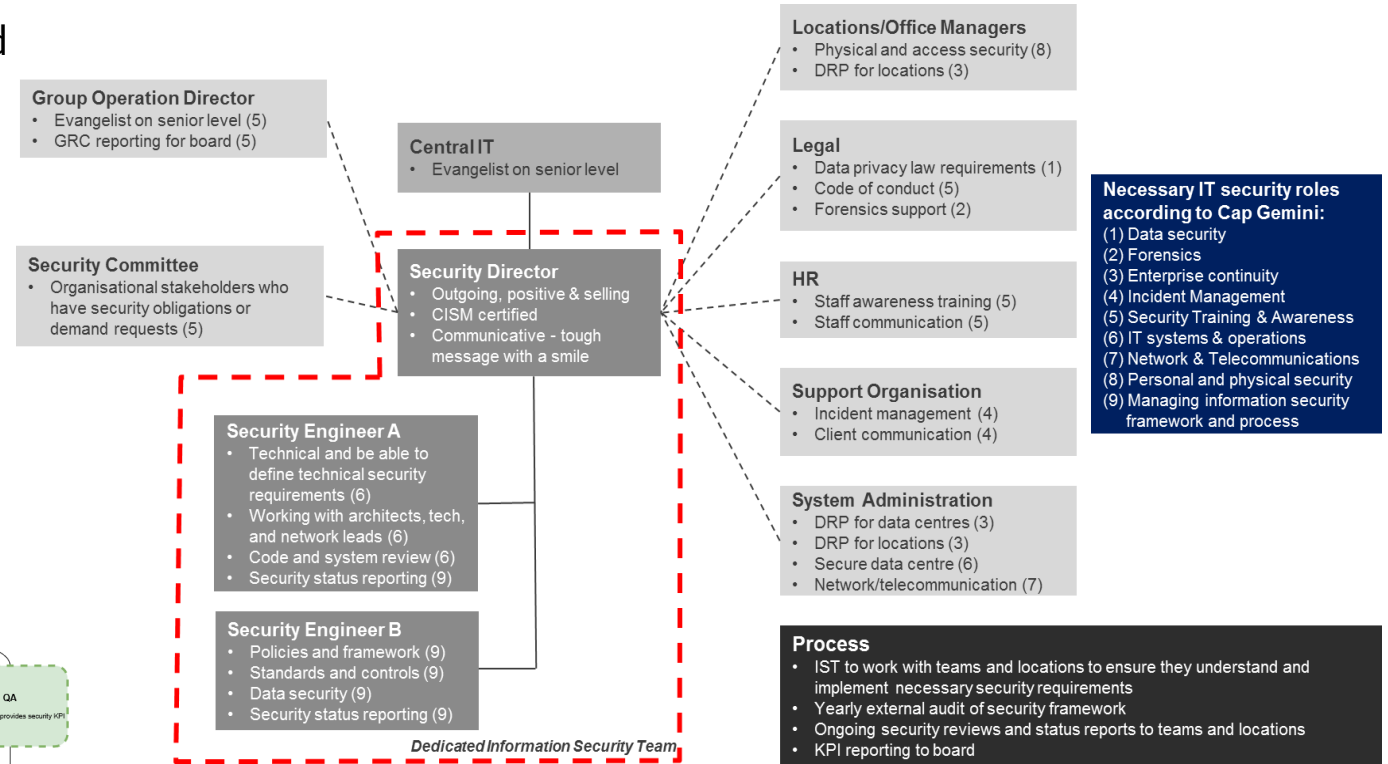
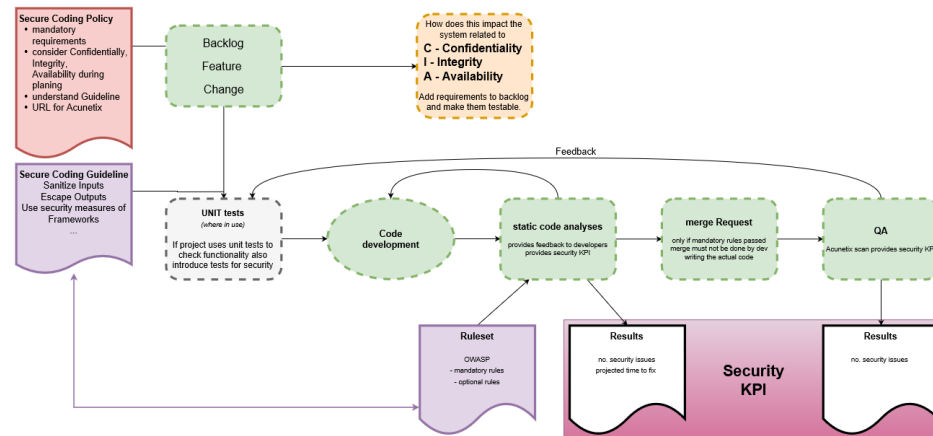
*Communicate, communicate, communicate*



# Dedicated Information Security Team

## Setup of Information Security Team

- Independent – can escalate to CEO and Board
- Policy framework
- Secure development guidelines
- System evaluation and guidance
- Open source scanning



# Dedicated Information Security Team

## Setup of Information Security Team

- Independent – can escalate
- Policy framework
- Secure development guide
- System evaluation and guidance
- Open source scanning

**Best practices for building resiliency**

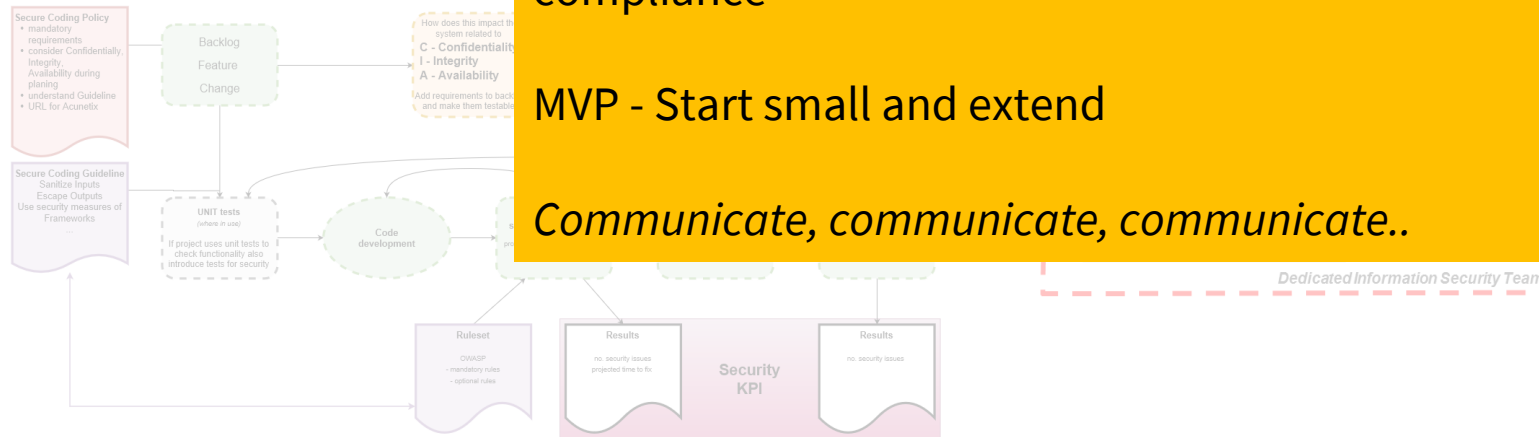
Include security as soon as possible

Train the trainer (eg. Tech Lead) concept

KPI's introduced by Policy used to measure compliance

MVP - Start small and extend

*Communicate, communicate, communicate..*



## Process

**Locations/Office Managers**

- Physical and access security (8)
- DRP for locations (3)

**Legal**

- Data privacy law requirements (1)
- Code of conduct (5)
- Forensics support (2)

**HR**

- Staff awareness training (5)
- Staff communication (5)

**Support Organisation**

- Incident management (4)
- Client communication (4)

**System Administration**

- DRP for data centres (3)
- DRP for locations (3)
- Secure data centre (6)
- Network/telecommunication (7)

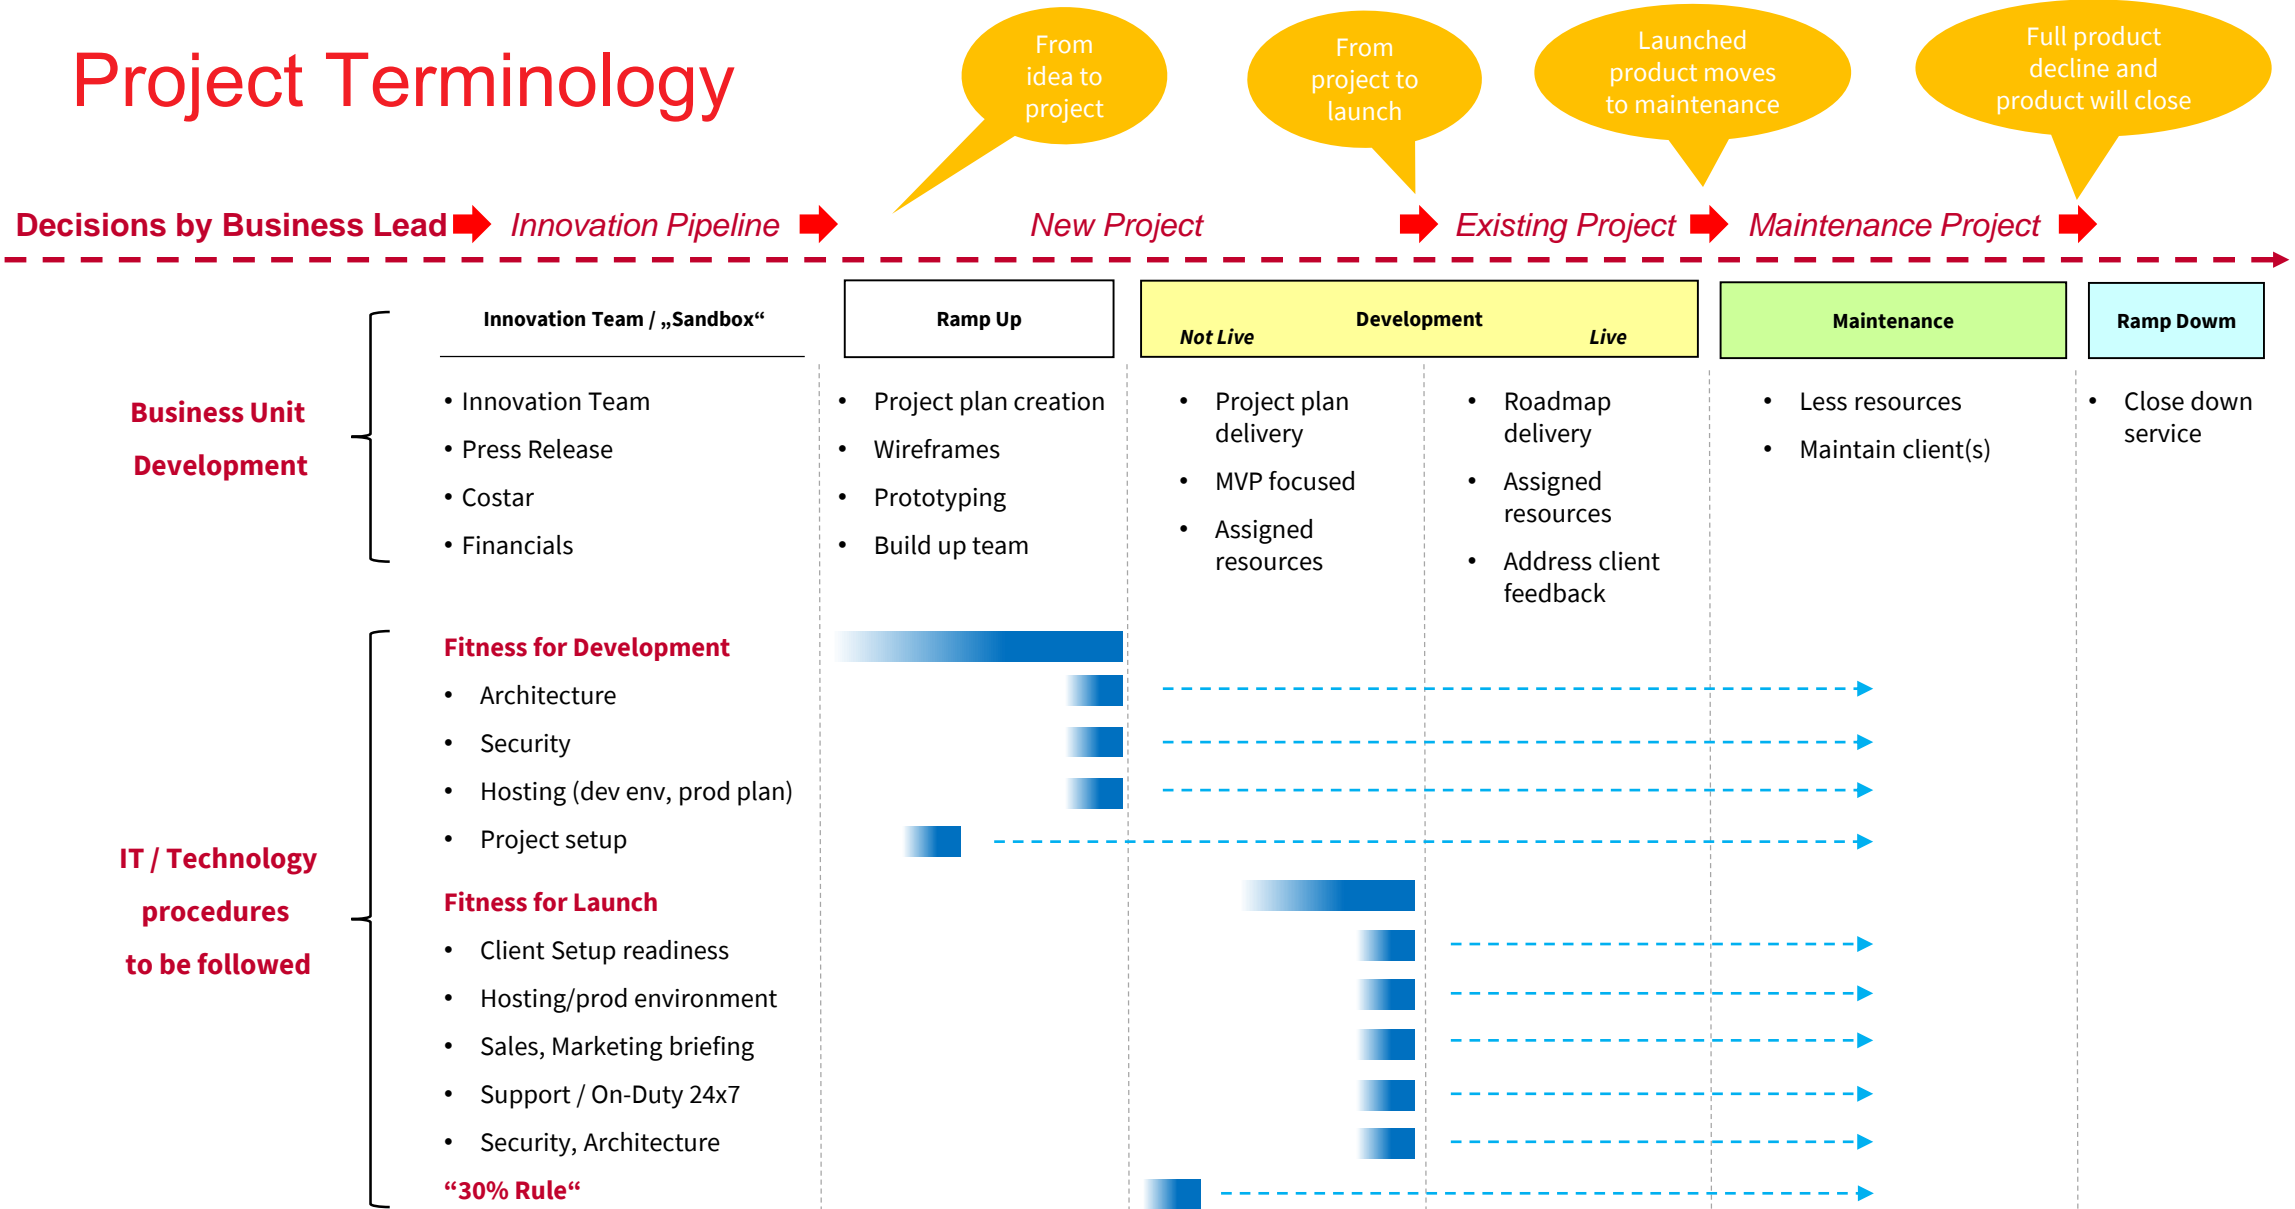
**Necessary IT security roles according to Cap Gemini:**

- (1) Data security
- (2) Forensics
- (3) Enterprise continuity
- (4) Incident Management
- (5) Security Training & Awareness
- (6) IT systems & operations
- (7) Network & Telecommunications
- (8) Personal and physical security
- (9) Managing information security framework and process

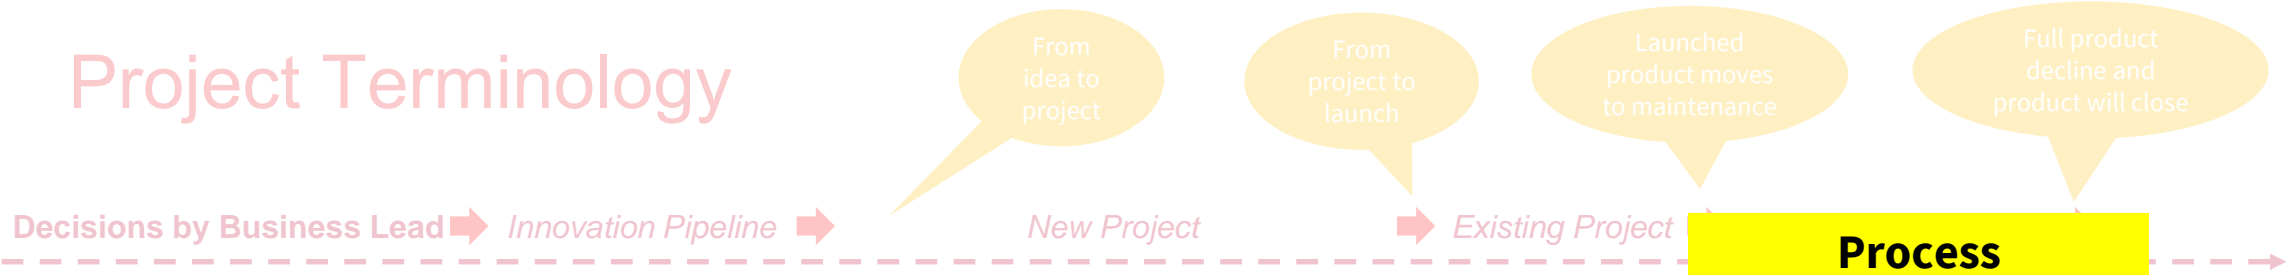
**Process**

- IST to work with teams and locations to ensure they understand and implement necessary security requirements
- Yearly external audit of security framework
- Ongoing security reviews and status reports to teams and locations
- KPI reporting to board

# Project Terminology



# Project Terminology



**Business Unit Development**

- Innovation
- Innovation
- Press Release
- Costar
- Financials

**Best practices for building resiliency**

Identify your key IT gate points

Do active project portfolio management so you know what is going on

Development teams in business unit – focusing on product features - tend to forget procedures, maintenance and stability so we have “30% Rule”

Iterate, remind and communicate procedures

Train the trainer (eg. Tech Lead) concept

**IT / Technology procedures to be followed**

- Fitness for Development**
  - Architecture
  - Security
  - Hosting (d)
  - Project se
- Fitness for Launch**
  - Client Set
  - Hosting/p
  - Sales, Mar
  - Support /
  - Security, Architecture

“30% Rule”


## Process

**Maintenance**

- Less resources
- Maintain client(s)

**Ramp Down**

- Close down service



# Best practices building resiliency - Wrap Up

## How to ensure all areas and systems are included

### *Technology*

- Definition of simple and clear technology architecture rules
- Ensure you use several hosting locations in a combined build up; either fully AWS or combination of on premise hosting with AWS
- Separate your service deliveries in logical pieces; frontend/backend, sub system clusters
- Understand your vendor dependencies

### *Cross IT challenge*

- Grey zone between development, system administration, hosting – clear DevOps topic
- Building resiliency is not only a technical topic; it's also about people processes and physical footprint

### *Governance*

- Enforce key IT processes; eg. Fitness for Development, Fitness for Launch and “30% Rule”
- Active project portfolio management with key IT processes as mandatory gate points
- Strict defined maintenance, support and incident processes

### *Continuous improvement*

- What went wrong – improve !!
- Peer review
- Postmortems

# Best practices building resiliency - Wrap Up

How to ensure all areas and systems are included

## Technology

- Definition of simple architecture
- Ensure you use several availability zones, either fully AWS or on-premise
- Separate your services into frontend/backend, stateless/stateful
- Understand your vendor dependencies

***Looks bureaucratic but it doesn't feel so***  
  
 (Comment from a Sportradar Tech Lead)

Business for Development, "le"  
 ment with key IT processes as  
 port and incident processes

## Cross IT challenge

- Grey zone between development, system administration, hosting – clear DevOps topic
- Building resiliency is not only a technical topic; it's also about people processes and physical footprint

## Continuous improvement

- What went wrong – improve !!
- Peer review
- Postmortems

# Sportradar

sportradar

Pablo Jensen, CTO

[p.jensen@sportradar.com](mailto:p.jensen@sportradar.com)

"y": 21.64

[www.linkedin.com/in/pablojensen](https://www.linkedin.com/in/pablojensen)



"x": 6.45

Thank you. |