BBC Studies of ST 2110 and NMOS for an On-Premise Cloud

Peter Brightwell
BBC

BBC R&D

IP SHOWCASE THEATRE AT IBC2019 : 13–17 SEPT 2019

Overview

- Why cloud? Why on-premise?
- BBC Cloud-Fit Production project
- Our R&D cloud environment
- Experiments
- Challenges for 2110
- Relationship with NMOS
- What next?

BBC R&D

IP SHOWCASE THEATRE AT IBC2019 : 13–17 SEPT 2019 2

Why migrate to IP?

- Future proofing for new content formats
- Support new ways of working
- Enable dynamic assignment of resources
- Scale facilities more easily
- Support multi-site and multi-tenant workflows
- Benefit from adopting COTS hardware on generic platform

A survey at NTS 2019…

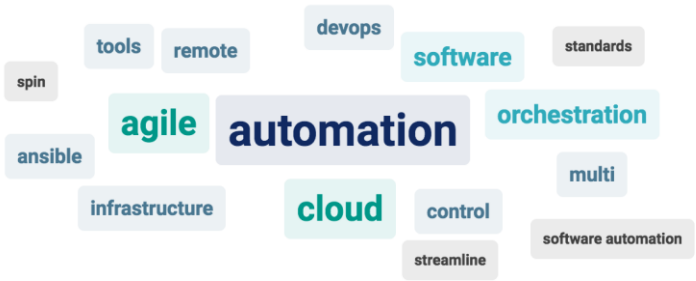

What innovative way of working would bring the biggest benefit to your organisation?

BBC R&D
IP SHOWCASE THEATRE AT IBC2019 : 13–17 SEPT 2019

Public and on-premise

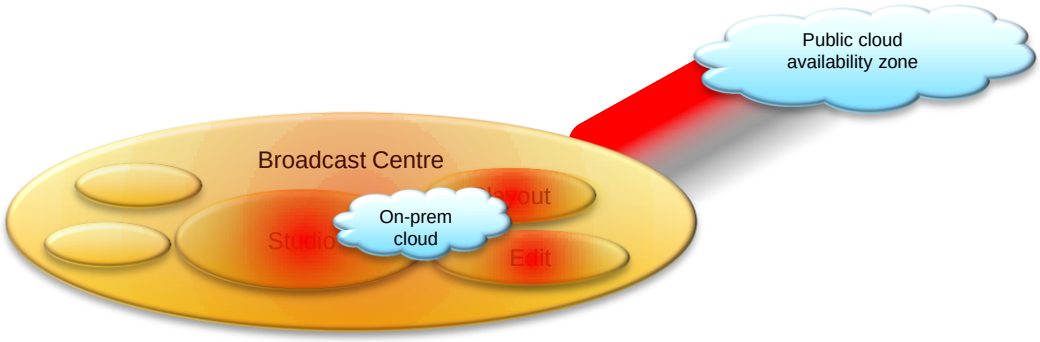

BBC R&D

On-premise cloud

Current cost models for public cloud not realistic for many productions

Potential to minimise power consumption

Need for specialisation

- Low latencies
- Fast access to huge amounts of local content
- High bandwidths
- “Non-typical” equipment

You know where your data is

…in practice both on-premise and public cloud will be likely

BBC R&D

BBC Cloud-Fit Production

Highly parallel approach for use with on-premise or public clouds

Computing, network and storage are abstracted as resources, deployed through services:

- Distributed
- User-defined
- Highly-available
- API-based

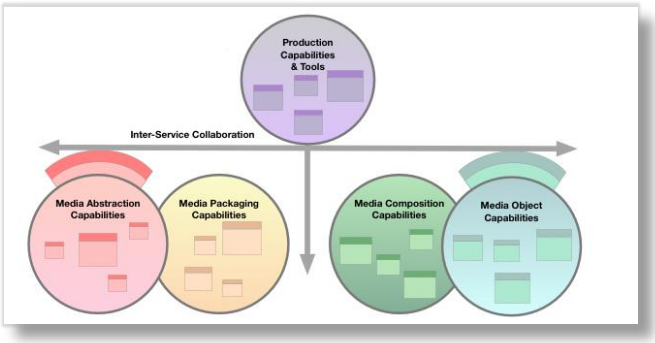

www.bbc.co.uk/rd/projects/cloud-fit-production

BBC R&D

IP SHOWCASE THEATRE AT IBC2019 : 13–17 SEPT 2019

Media Object Store (Squirrel)

- Parallelised storage of streams
- Dividing into objects and load balancing allow scalability
- Database allows easy access
- S3 API supported by most cloud frameworks

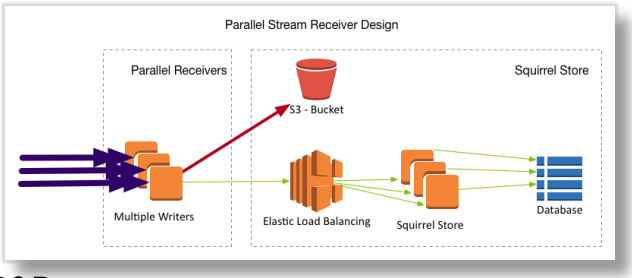

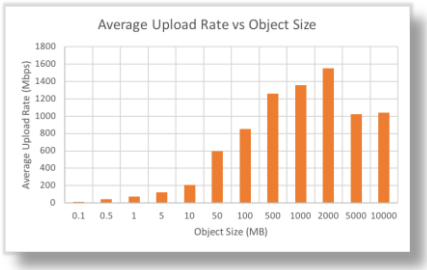

BBC R&D

IP SHOWCASE THEATRE AT IBC2019 : 13–17 SEPT 2019

Media File Ingest (Magpie)

- Chunker breaks up source file
- Reader queues chunks into Media Store
- Grain Headers provide IDs and timestamps

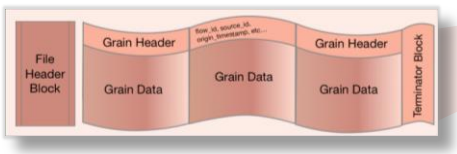

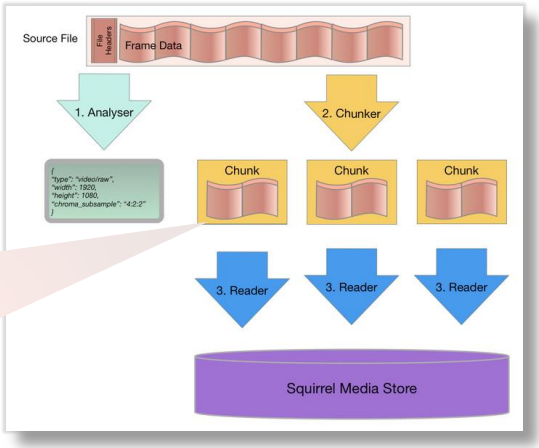

BBC R&D

IP SHOWCASE THEATRE AT IBC2019 : 13–17 SEPT 2019

Stream Packaging Service (Otter)

- Fetcher gets objects from Media Store
- Dispatcher turns them into streams
  - First work has used FFMpeg
  - 2110-based Dispatcher in development

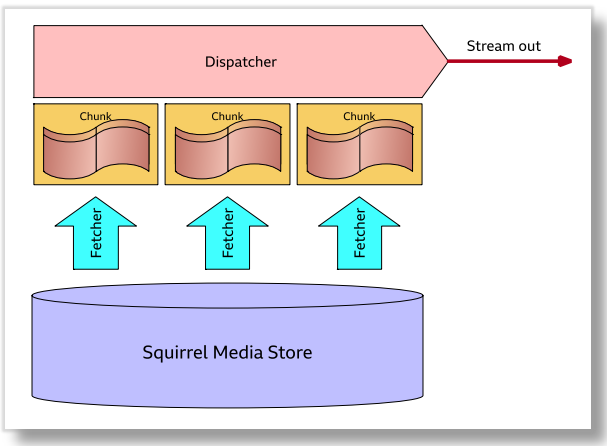

BBC R&D

IP SHOWCASE THEATRE AT IBC2019 : 13–17 SEPT 2019 10

IP SHOWCASE THEATRE

BBC R&D's prototype cloud

Allows us to deploy wherever is required
- Isolated networking and compute
- Multi-tenanted
- Containerised applications

Automated and repeatable
- Wipe and rebuild in minutes

Learn from data centre architectures
- Leaf-and-spine
- Multi-chassis LAG
- BGP-EVPN / VXLAN overlay

Avoid lock-in
- Open source tooling
- Non-proprietary networking
- COTS hardware

BBC R&D

IP SHOWCASE THEATRE

openstack.

OpenStack

Open source framework
often used to self-provide clouds

Created in 2010 by NASA and Rackspace

Used by: CERN, telcos, Adobe, Blizzard, Oath…

APIs provide infrastructure as a service ("IaaS")

Not (necessarily) an all-in-one solution
- Can be used with third party components, e.g. Ceph object storage, vmware hypervisor

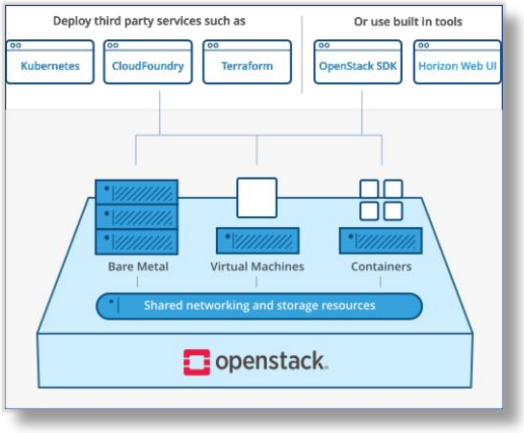

BBC R&D

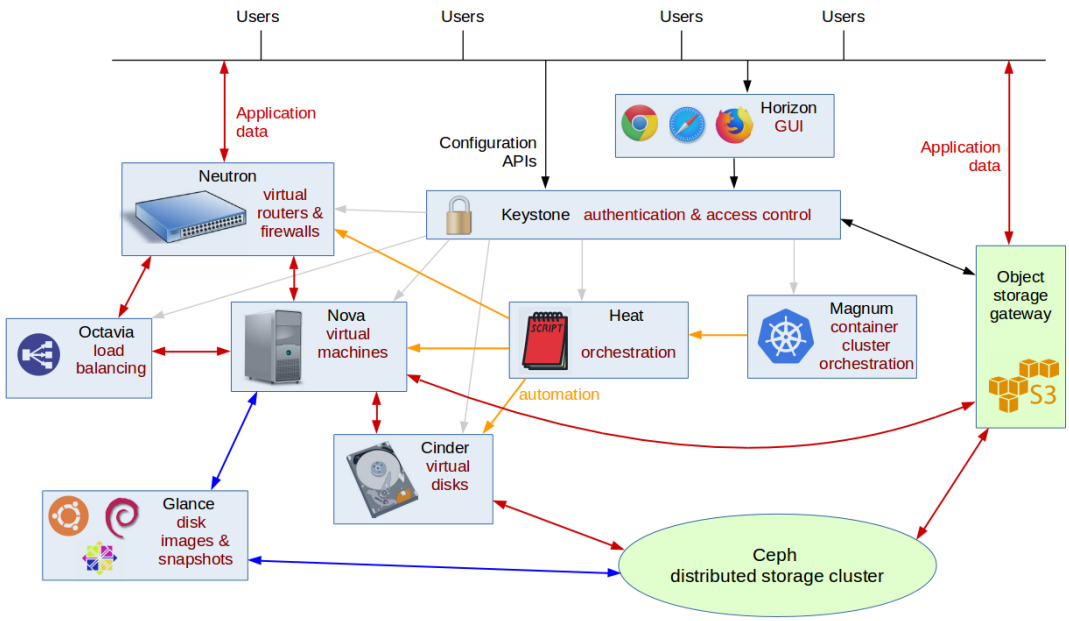

BBC R&D

IP SHOWCASE THEATRE AT IBC2019 : 13–17 SEPT 2019

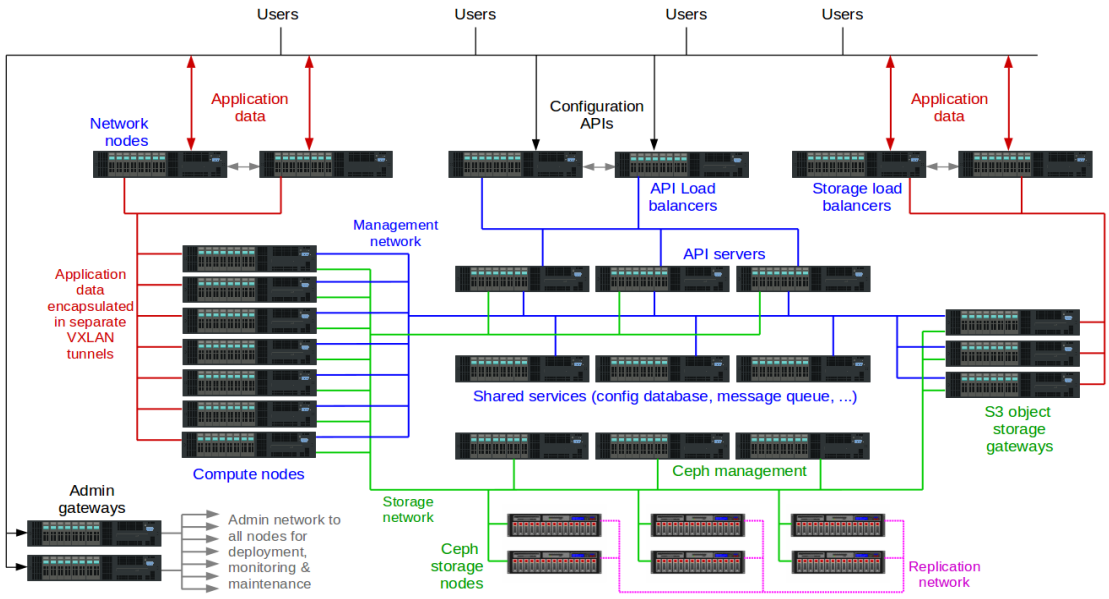

BBC R&D

IP SHOWCASE THEATRE AT IBC2019 : 13–17 SEPT 2019

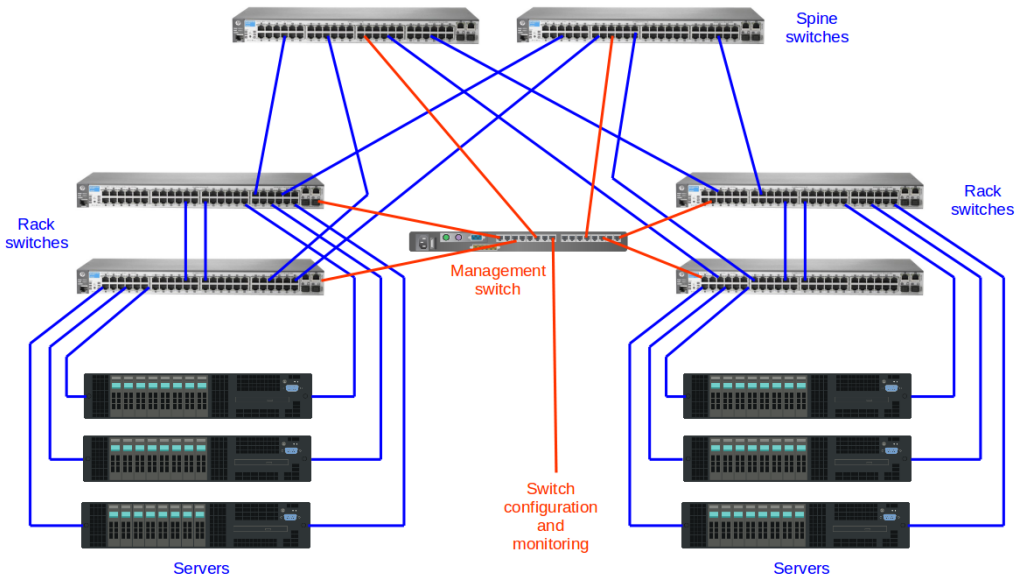

BBC R&D

IP SHOWCASE THEATRE AT IBC2019 : 13–17 SEPT 2019

IP SHOWCASE THEATRE

Automation with Ansible

SERVERS

Configure BIOS options e.g. boot sequence
Partition disks and create filesystems
Configure network settings and proxies
Configure NTP servers
Set up ***apt*** and ***pip*** package repositories
Install CA root certificates
Create user accounts
Install SSH public keys
Install required software packages
Create or modify software config files
Test network connectivity

SWITCHES

Configure layer 2 VLANs
Configure layer 3 VRFs
Configure access and trunk ports
Create aggregated port-channels, vPCs
Create layer 3 routed links between switches (BGP+EIGRP)
Create floating virtual IPs for transparent failover
Configure multicast and RPs
Set up DHCP relays

OPENSTACK HOSTS

Apply security hardening
Create containers and their network bridges
Install applications and libraries
Configure service accounts, kernel options, etc.
Create customised config files
Supply scripts and tools for maintenance and upgrades

BBC R&D

IP SHOWCASE THEATRE AT IBC2019 : 13–17 SEPT 2019 16

Salford

London

BBC R&D

Experiments with our on-premise cloud

- Uncompressed capture and stream through BBC's IP Studio network
- Gateway with kernel bypass (netmap)
- Local Ceph cluster with 30 nodes
- 1080/i25 video frames

www.bbc.co.uk/rd/blog/
2019-03-live-video-ingest-aws-openstack-cloud

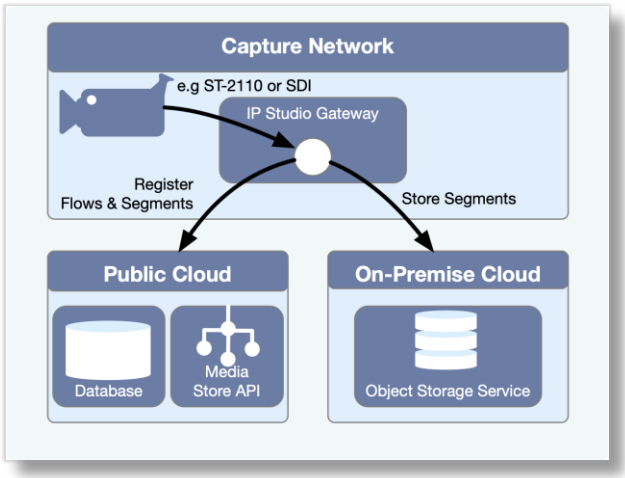

BBC R&D

IP SHOWCASE THEATRE

Ceph Dashboard

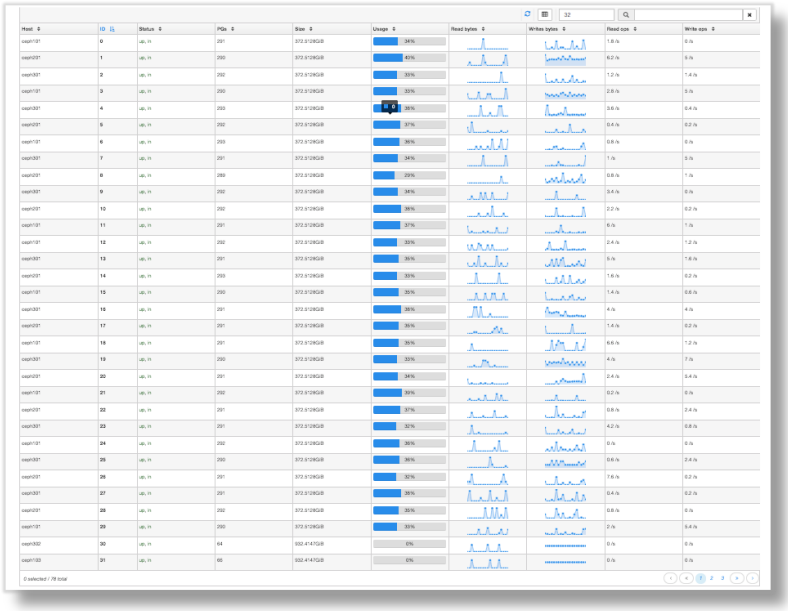

BBC R&D

IP SHOWCASE THEATRE

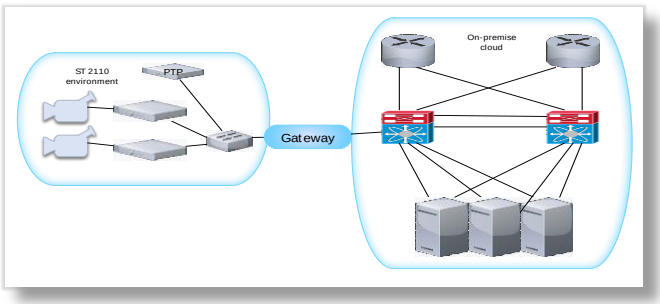

We are here…

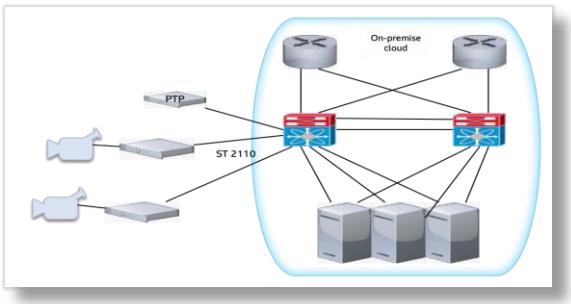

…should we, can we, get here?

BBC R&D

Challenges observed with virtualising ST 2110 endpoints

- Available compute performance varies over time and with host loading
- Virtualised NIC performance too poor to meet -21 constraints
- Directly attached NIC performs better but lacks flexibility
- Bridged network provides a compromise by did not work well with VXLAN
- SR-IOV differences between NICs affects virtual function performance

Effect is to make 2110-21 compliant working on virtualised platforms difficult

BBC R&D

Improving the performance

Use SR-IOV

Bypass the kernel (netmap, DPKD)

Off-load functionality onto the NIC

- Packet pacing
- Kernel bypass
- Packet aggregation
- VXLAN overlay

These come at the cost of increasing complexity and may restrict choice of vendors

BBC R&D

Further challenges

Broadcast live IP combines several requirements that aren't typical in many cloud scenarios

- IGMP multicast
- PTP based timing
- 2022-7 resilience patterns
- Minimal latency
- Dedicated resourcing

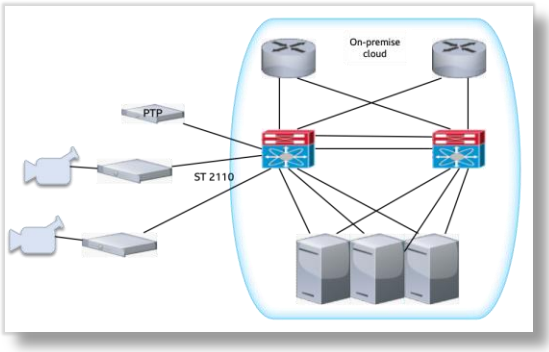

BBC R&D

What about NMOS?

NMOS could be confined to the ST 2110 environment…

- Gateway is a Node

…or could extend into the cloud

- Which would need an IS-04 Registry
- This could be in switch or compute

Within the cloud, other service discovery frameworks may become more relevant

In any case, hooking into wider orchestration will become important

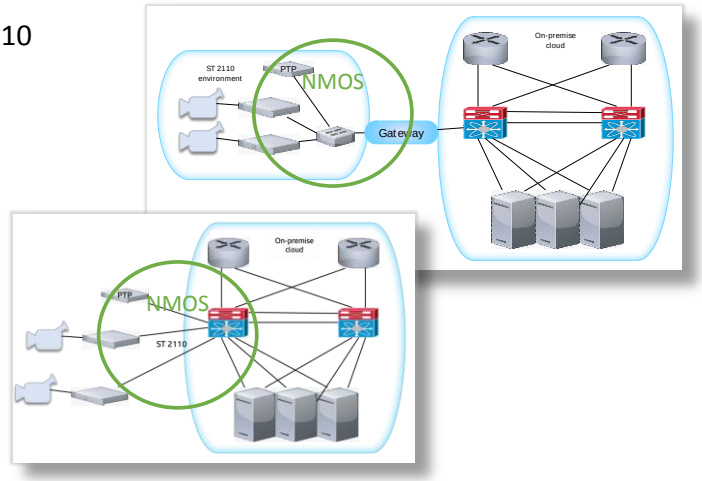

BBC R&D

IP SHOWCASE THEATRE

What are we doing next?

- Building the London cloud
- Learning from vendor deployment frameworks
- Integrating aspects of these into our OpenStack cloud
- 2110 output with Stream Packaging Service
- Building and trialling more cloud-fit components
- Considering the service discovery question
- (With EBU) Broadcaster requirements for automation and orchestration

BBC R&D

IP SHOWCASE THEATRE

Thank you

Peter Brightwell, BBC
peter.brightwell@bbc.co.uk
www.bbc.co.uk/rd

Thank you to our Media Partners

SVG SPORTS VIDEO GROUP
SVG SPORTS VIDEO GROUP EUROPE
BROADCAST THE BRIDGE Connecting IT to Broadcast
TVB EUROPE Intelligence for the media & entertainment industry

BBC R&D