AI in networks

Dr Simon Thompson
Principal Researcher in Customer Experience
and AI Readiness

The challenge

Next Generation Network platform – built around orchestration

We are investing in our platforms to give the desired agility and transparency

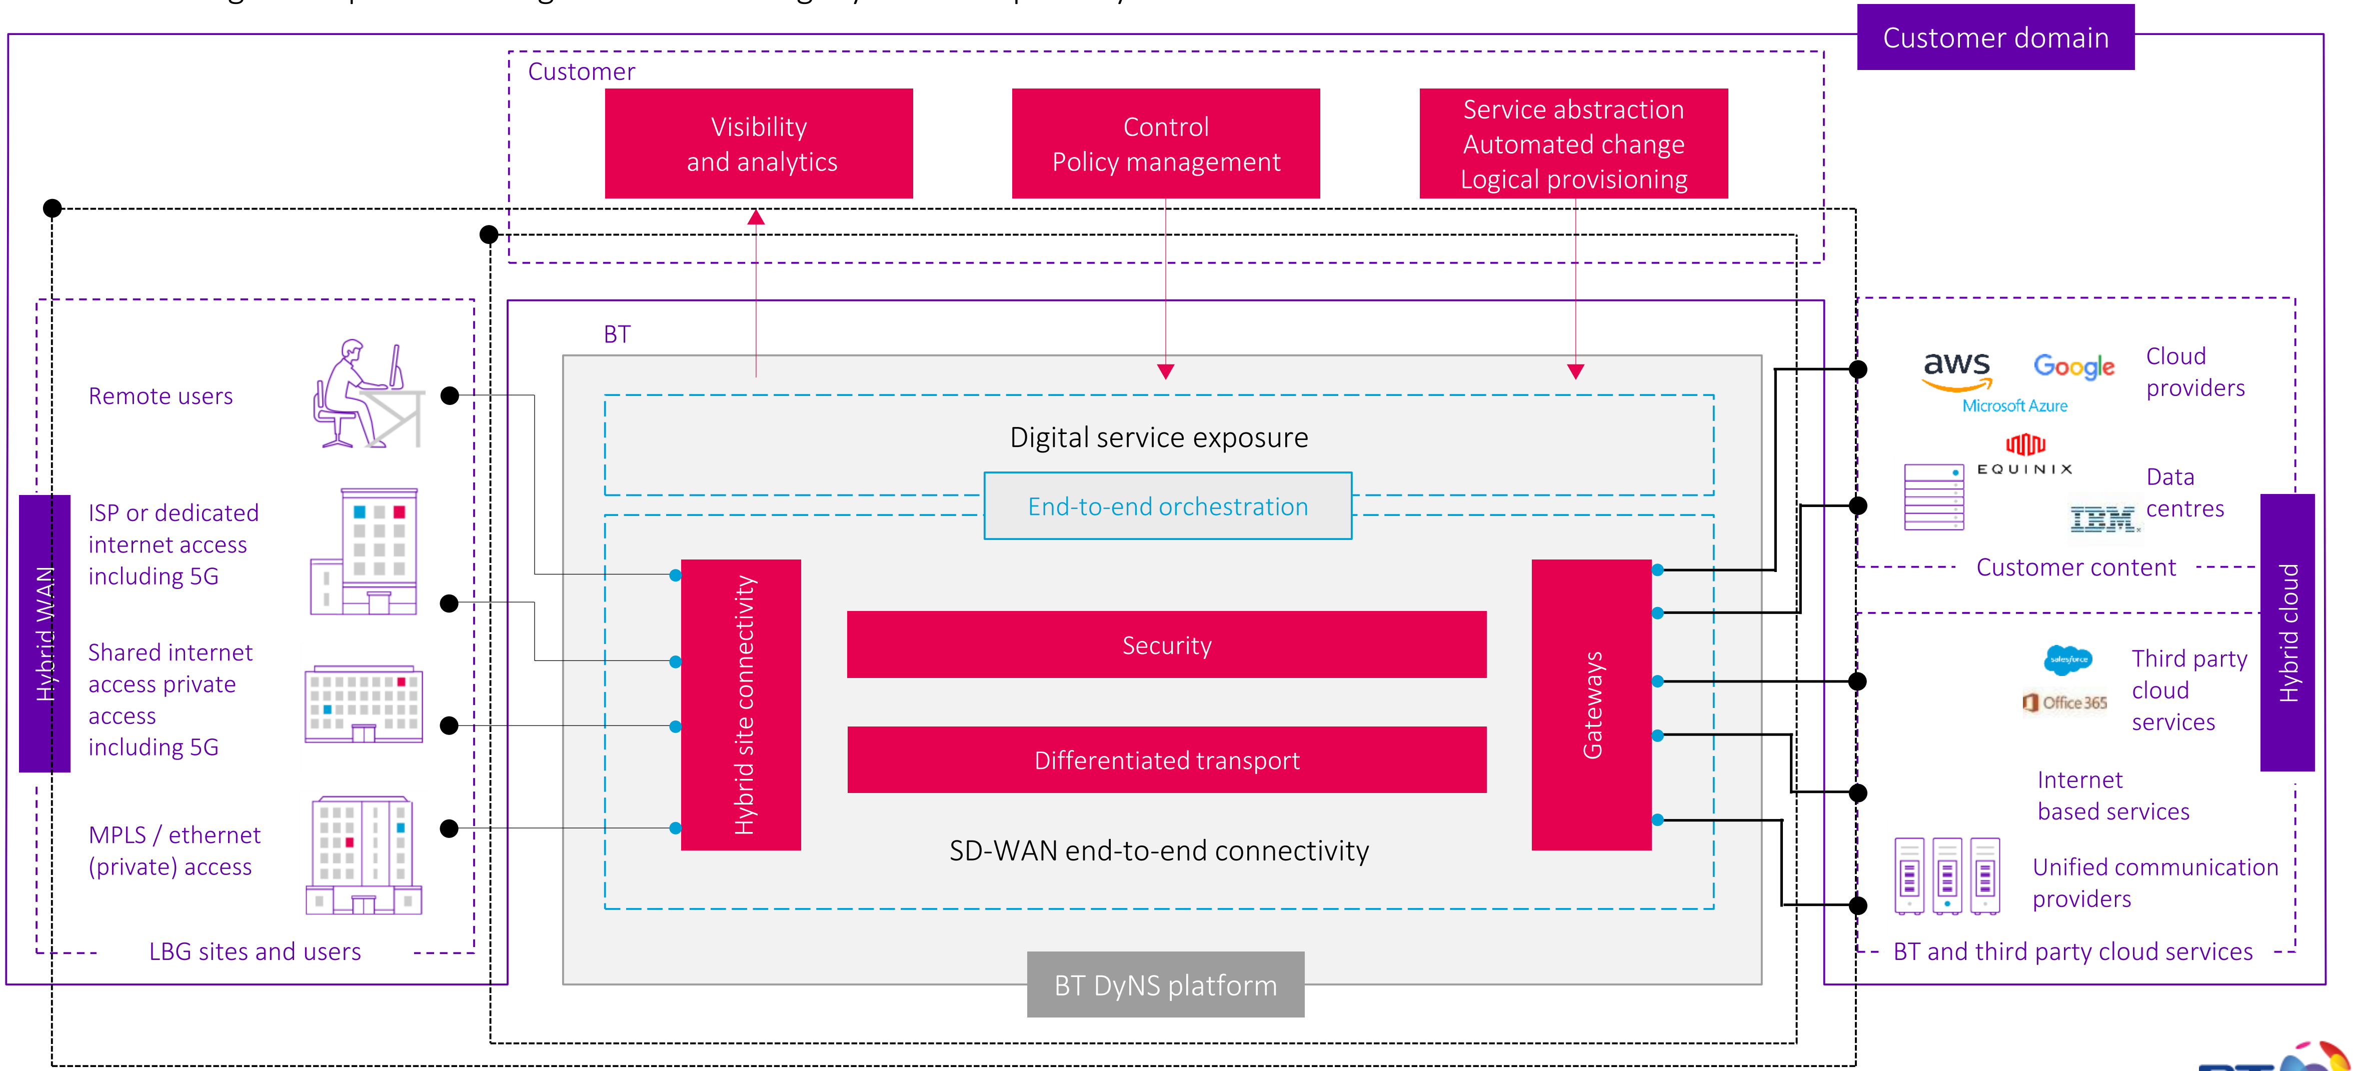

Proposition

People and processes

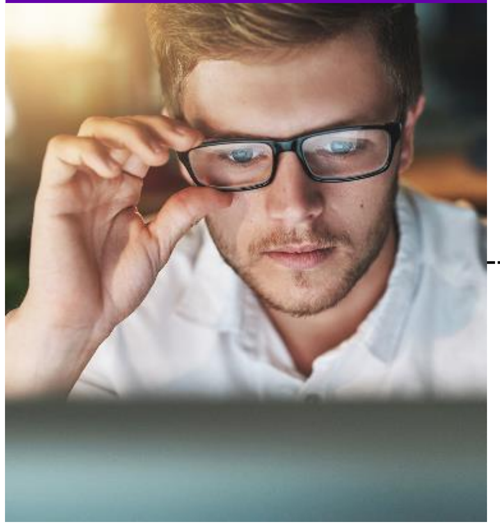

Network provider

Customer

End users

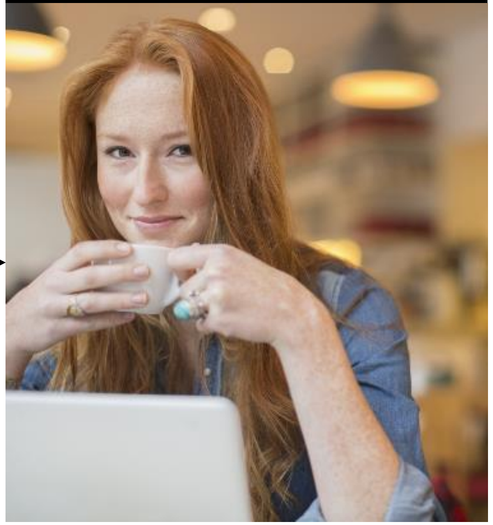

Utilisation
Knowledge
Flexibility

Control
Visibility
Responsiveness

Proposition

Network AI

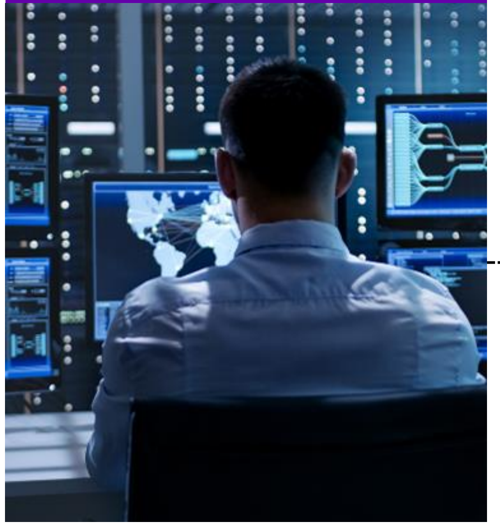

Network provider

Customer

End users

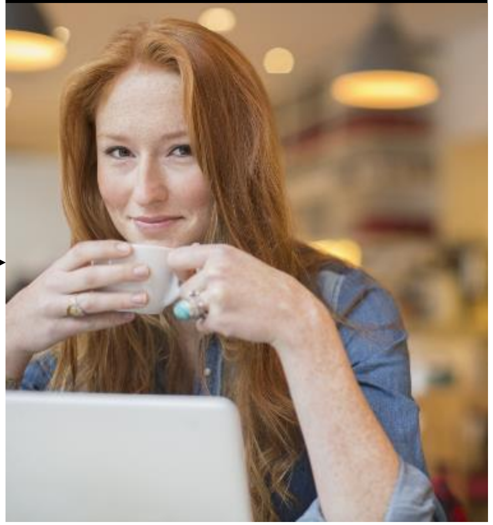

Utilisation
Knowledge
Flexibility

Control
Visibility
Responsiveness

Machine learning and closed-loop control in BT's programmable network

With the deployment of network telemetry along with model-based service creation closed-loop automation becomes possible

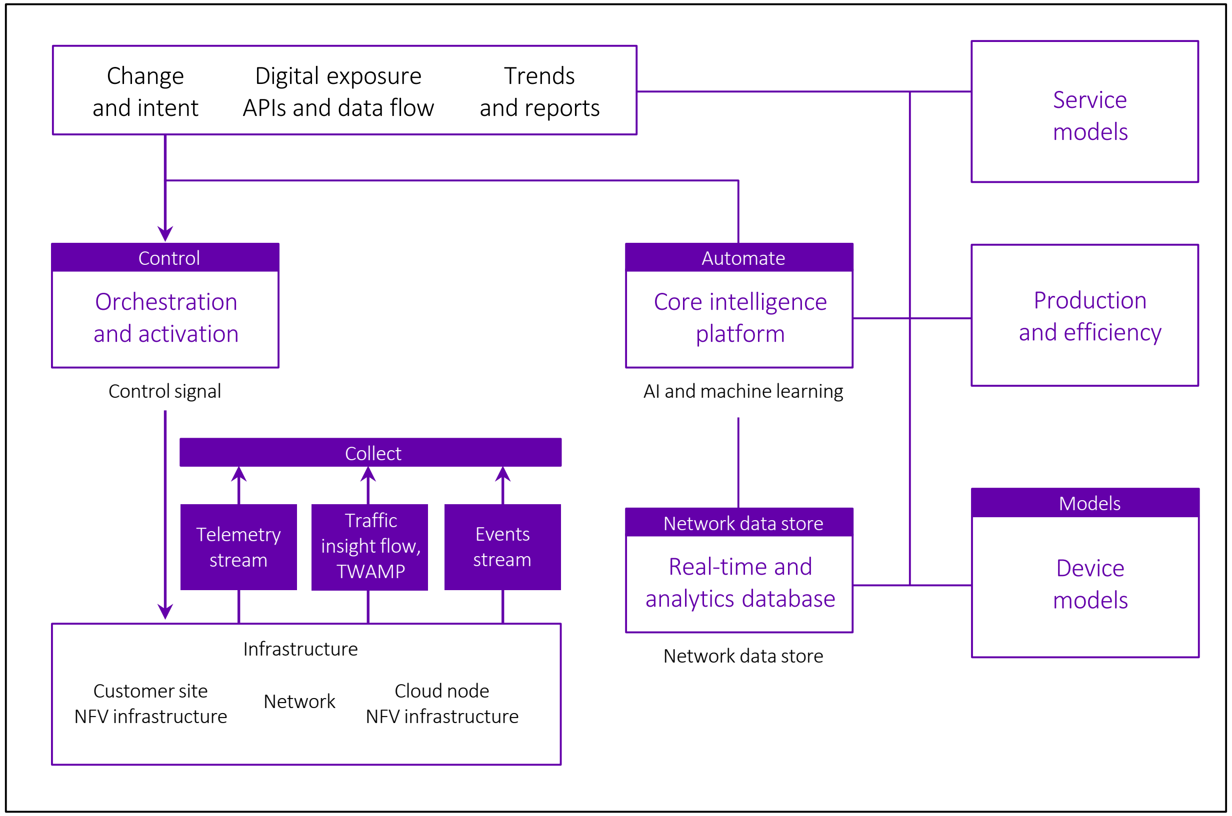

Collect

- Automated collection of network data into a real-time and analytics database
- Network statistics, flow, TWAMP, etc.
- Streamed model driven telemetry using YANG models.

Correlate

- Using baseline metrics correlate network changes / performance / problems
- Derive policy from network norms.

Automate (AI and ML)

- Detect anomalies and take pre-emptive action rather than relying on resilience or restoration
- Efficient, fair, and pragmatic resource allocation.

Control

- Use automation to drive service assurance, work-around failure scenarios
- Improve capacity planning
- Deal with the micro-outages
- Extended API's enabling abstraction of network functions.

Things of interest may only last a few seconds

- 100m final of 2012 Olympics lasted less than 10 seconds – potentially leading to very short term surge in demand for our network.
- Current knowledge of network performance gives us performance reports every few minutes – we need to know about performance every second and in real time to allow us to act and optimise service.

Telemetry will enable us to know far more about how the network is performing on a second by second basis – enabling us to optimise performance for our customers

Network AI in BT so far

Big data (2013 style)

Ad hoc investigation and exploration

DSLAM naughty list for supplier management

Analysis and report for monitoring

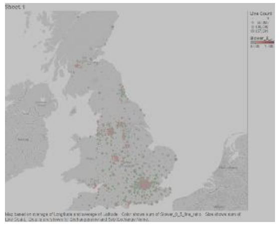

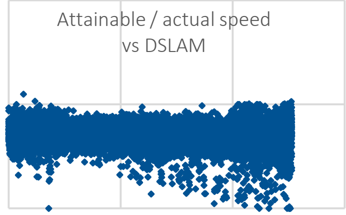

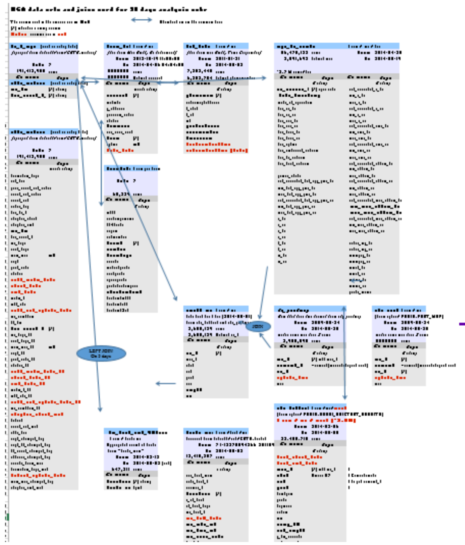

V. large joins to create Single truth per line

14 data assets describing different aspects of access network estate and performance

5.29.117.6	ECI_BFocusV2FubR_2.02.01	2701-100589-005	ECI	< 30%	1-1.5k	1-1.5k	0-30Mbps	1	12
6.3.9.63-plus.tm	ECI_BFocusV2FubR_2.02.01	2701-100588-005	ECI	70-89%	1.5+k	1.5+k	0-30Mbps	1	8
4.7.5.1.83.8.130.1.17	Huawei_EchoLife HG612_V100R001C01B028SP10TELWEB	Home Hub 40 Type A	Huawei	> 90%	0-1k	0-1k	60+Mbps	1	6
4.7.5.1.83.8.130.1.5.1.1	ECI_BFocusV2FubR_2.02.01	Home Hub 40 Type A	ECI	70-89%	1-1.5k	0-1k	60+Mbps	1	6
4.7.5.1.83.8.173.1.6	ECI_ECLVL05_1.0.5.b89b	Home Hub 50 Type A	ECI	30-49%	0-1k	1.5+k	0-30Mbps	1	6
V100R001C01B036SP05_L_B	Huawei_EchoLife HG612_V100R001C01B028SP10TELWEB	Home Hub 30 Type B	Huawei	< 30%	1-1.5k	1.5+k	0-30Mbps	1	6
4.7.5.1.83.8.130.1.5.1.1	Huawei_EchoLife HG612_V100R001C01B030SP06	Home Hub 40 Type A	Huawei	50-69%	1.5+k	0-1k	60+Mbps	2	11
4.7.5.1.83.8.173.1.4		Home Hub 50 Type A	ECI	70-89%	1.5+k	1-1.5k	0-30Mbps	1	5
4.7.5.1.83.8.173.1.4	ECI_BFocusV2FubR_2.02.01	Home Hub 50 Type A	ECI	70-89%	1.5+k	1.5+k	0-30Mbps	1	5

Speed

0	0	0	0	3
0	0	6	3	0
0	0	0	0	0
0	0	0	0	0

Line length

Modem type

Bucketed cube of performance vs Line / dp / dslam / exchange / supplier / Type

Bayesian AI

The cumulative errors in our records and the natural variation in performance (e.g. through between pair variance) make it difficult to apply traditional AI techniques to the problems we face and the complex topological relationships are unsuited to modern methods such as deep learning.

Here we present a Bayesian probabilistic approach to modelling the access network which allows us to overcome the noise and uncertainty in our data to extract knowledge and value.

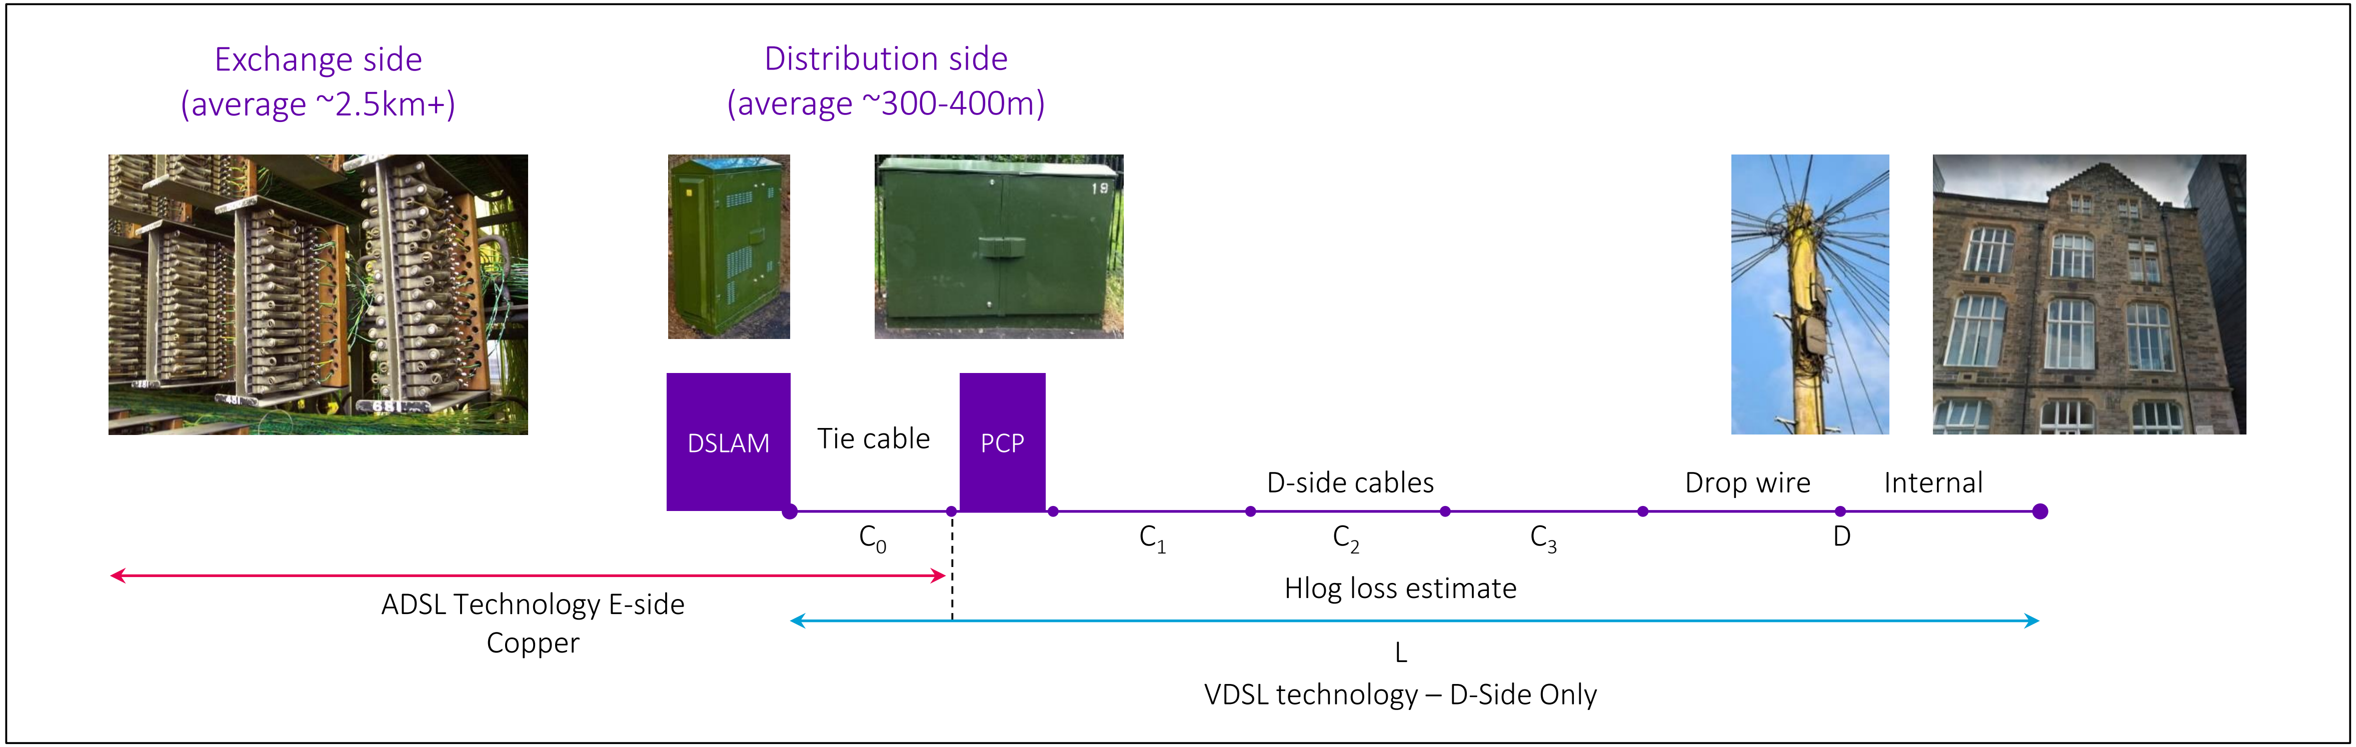

Anomaly detection for predictive network analytics

For each network element (10,000s)

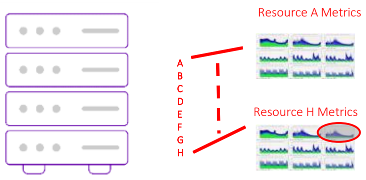

1. Process (Big data)

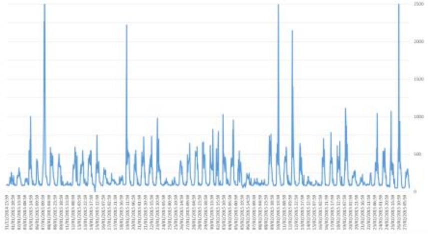

- For every resource / metric.
- Clean and snap to regular intervals.

2. Feature extraction

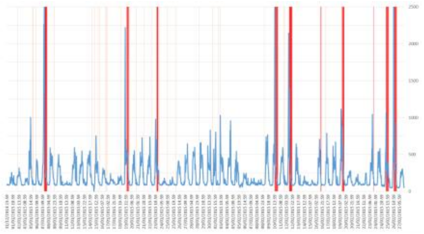

- Run anomaly detector suite.
- Data reduction.
- Produce patterns for diagnostics.
- Several algorithms in parallel (Cortical learning, PELT, SAX).

3. Prediction

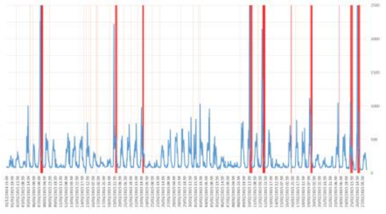

- Data volume greatly reduced.
- Only process anomalies.
- Continuous M / learning desirable.

4. Apply

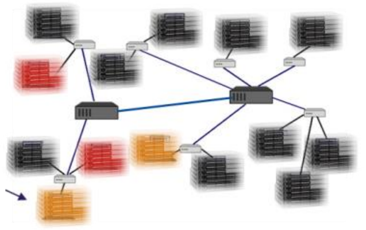

- Record probabilities for each outcome.
- Flag warnings and visualise probabilities (in space and time).
- Use Bayesian Networks.

Using Deep Learning to understand network equipment

4 We train this bit on our data, this gives us the power of the pre-trained neural network, but only requires a few hundred or so training images for good results.

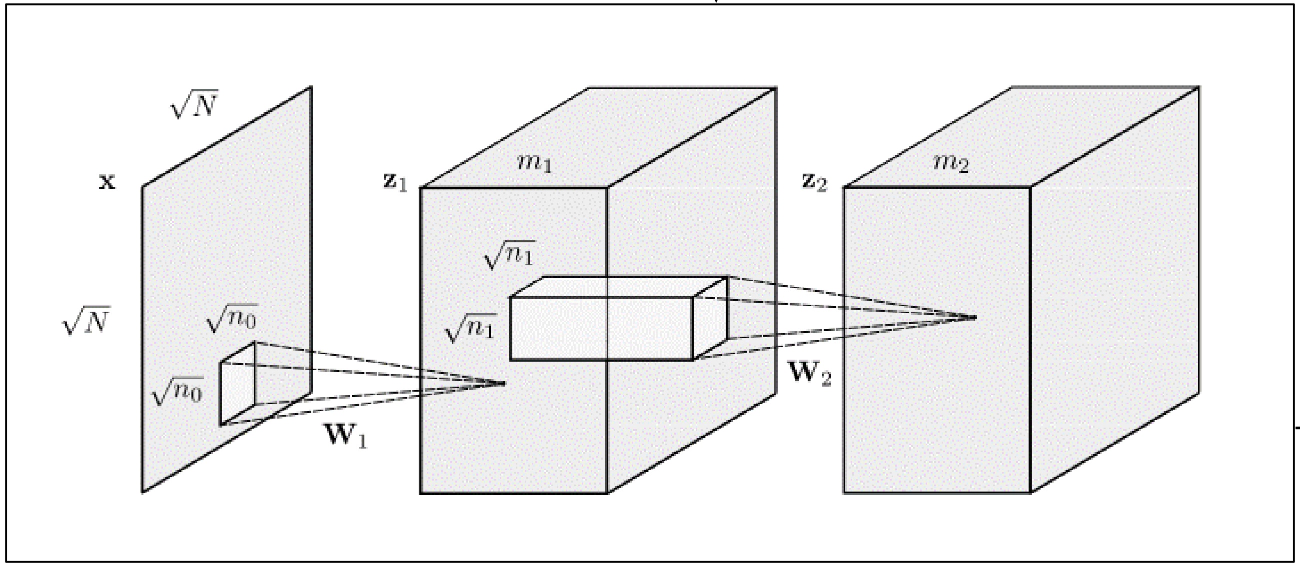

1 Deep model trained on generic image data sets (crowd sourced labels)

2 Now we can use the deep network to make inferences about images captured in the field

- How much capacity is available
- What kinds of technology
- State of repair

3 Small expert labelled training set of images from surveys.

Creating value with AI in the network – ‘enabling a zero touch network’

For 400 premises we approximate that we require initially 100 pieces of equipment. For each piece of equipment there are say 100 possible locations. That’s 100100 possible combinations in which the equipment can be equipment can be located (or a googol of googols).

Evaluate a solution every millisecond (3.153610 ms per year) then it will take 1.4199 years to evaluate all possible solutions.

Determining the best solution deterministically is not possible.

To put in perspective, the lifetime of the universe would run as follows:

The sun dies before 10^{10} years
At 10^{14} years, star formation ceases
At 10^{15} years, planets are flung from their orbits
At 10^{19} years, galaxies no longer exist
At 10^{40} years, all protons have decayed

Creating value with AI in the network – ‘enabling a zero touch network’

Network model data

For example:

- Network nodes / links
- Equipment / cables connections
- Demand points
- Distribution points
- Network costs
- Bandwidth / service requirements

Network design rules

For example:

- Equipment / cable capacity
- Equipment location rules
- Node / link characteristics
- Power budgets
- Equipment configurations
- Design parameters
- Network topology

NetDesign

- AI-Heuristic search e.g.
  - Guided local search
  - Genetic algorithms
  - SA’s
- Graph theory
  - Routing
  - Trees
- Multi layered network modelling

Optimised network designs

- Node / link locations
- Network topologies
- Bill of materials
- Schematics
- Connectivity

Industry recognition

Using an AI model to identify collaborative solutions across a wide set of stakeholders

We can use this AI model to generate new collaborative solutions across a wide set of stakeholders

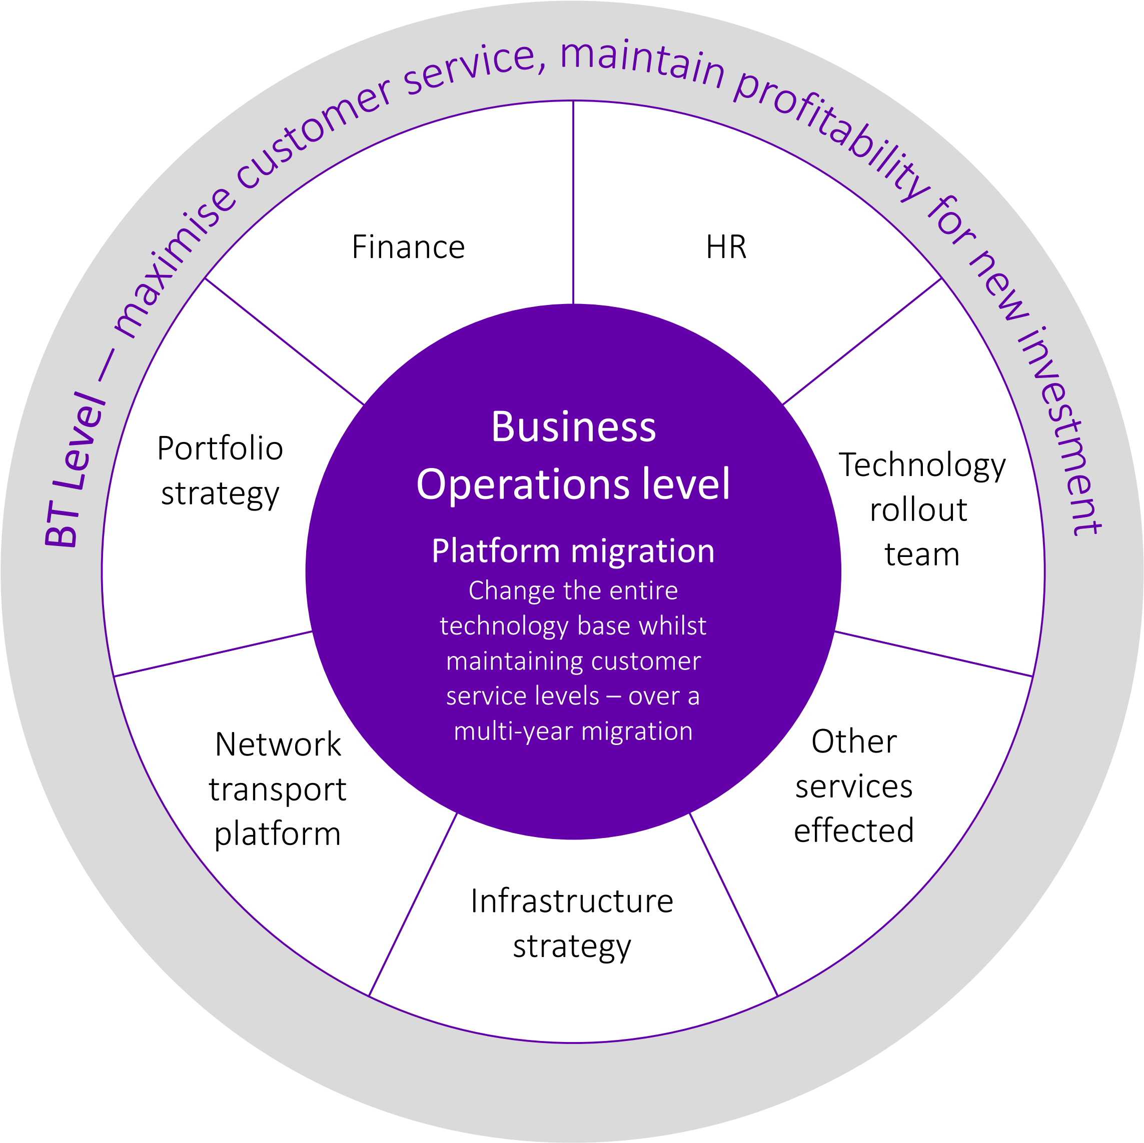

Again – vast combinatorial search, AI mediates with humans

Complete detailed work plan (17m steps)

Multi-year costs, system size, etc…

A1141	unit1141	B1141	TUN	SC	HW	clstr	Mark	lineTest	Aused	Afree	Afitted	I2used	I2free	I2fitted	I30used	I30free	I30fitted
AB/LN/UD	AGV/UB	AGV/UB/5	3	6	48	1	3	No Data	0(1)	1(0)	1	0(1)	1(0)	1	216(268)	52(0)	270
AB/LN/UD	AGV/UB	AGV/UB/3	5	6	38	1	2	Master/1	579(1920)	1341(0)	1920	8(16)	8(0)	16	0(0)	0(0)	0
AB/LN/UD	AGV/UB	AGV/UB/4	5	7	39	1	2	Master/1	272(1024)	752(0)	1024	65(56)	31(40)	96	0(0)	0(0)	0
AB/LN/UD	AGV/UB	AGV/UB/1	3	0	0	1	1	Master/1	1658(77)	198(1779)	1856	0(0)	1(1)	1	75(53)	15(37)	90
AB/LN/UD	AGV/UB	AGV/UB/2	3	1	1	1	1	Master/1	513(0)	1215(1728	1728	0(0)	1(1)	1	30(0)	60(90)	120
ACM/PG	CAM/UE	CAM/UE/1	5	2	35	8	1	Master/1	596(1856)	1260(0)	1856	0(1)	1(0)	1	0(0)	0(0)	0
ACM/PG	CAM/UE	CAM/UE/7	4	1	21	8	1	Master/1	549(451)	795(893)	1344	30(64)	34(0)	64	10(30)	20(0)	30
ACM/PG	CAM/UE	CAM/UE/1	5	0	24	8	1	Master/1	532(0)	748(1280)	1280	0(1)	1(0)	1	0(0)	0(0)	0
ACM/PG	CAM/UE	CAM/UE/6	4	0	20	8	1	Master/1	408(0)	744(1152)	1152	0(1)	1(0)	1	0(0)	30(30)	30
ACM/PG	CAM/UE	CAM/UE/1	3	0	15	8	1	Master/1	222(0)	34(256)	256	96(59)	64(101)	160	20(0)	40(60)	60
BD/2B	LBS/UG	LBS/UG/1	7	0	5	2	3	Master/1	1524(1792	268(0)	1792	0(1)	1(0)	1	14(14)	4(4)	30
BD/2B	LBS/UG	LBS/UG/2	7	1	6	2	3	Master/1	1465(1792	327(0)	1792	0(1)	1(0)	1	0(0)	0(0)	0
BD/2B	LBS/UG	LBS/UG/3	8	0	7	2	3	Master/1	624(764)	1424(1284	2048	0(1)	1(0)	1	0(0)	0(0)	0
BD/2B	LBS/UG	LBS/UG/4	8	1	8	2	3	Master/1	445(0)	1475(1920	1920	25(32)	7(0)	32	0(0)	0(0)	0
BD/2B	LBS/UG	LBS/UG/5	8	2	9	2	3	Master/1	290(0)	350(640)	640	59(49)	69(79)	128	0(0)	0(0)	0
BD/DA	HF/UF	HF/UF/6	11	3	11	3	3	No Data	0(1)	1(0)	1	0(1)	1(0)	1	0(0)	0(0)	0
BD/DA	HF/UF	HF/UF/1	10	2	6	3	2	Master/1	1558(1792	234(0)	1792	0(1)	1(0)	1	0(0)	0(0)	0
BD/DA	HF/UF	HF/UF/2	10	3	7	3	2	Master/1	1170(1408	238(0)	1408	0(1)	1(0)	1	0(0)	0(0)	0
BD/DA	HF/UF	HF/UF/5	11	2	10	3	2	Master/1	708(888)	1020(840)	1728	63(80)	17(0)	80	0(0)	0(0)	0
BD/DA	HF/UF	HF/UF/3	11	0	8	3	2	Master/1	653(0)	1011(1664	1664	48(28)	-16(4)	32	0(0)	0(0)	0
BG/DPL	CA/UB	CA/UB/4	7	2	34	2	2	Master/1	1622(2048	426(0)	2048	0(1)	1(0)	1	0(0)	0(0)	0
BG/DPL	CA/UB	CA/UB/3	6	4	33	2	2	Master/1	865(1920)	1055(0)	1920	30(32)	2(0)	32	0(0)	0(0)	0
BG/DPL	CA/UB	CA/UB/2	6	3	32	2	2	Master/1	824(621)	968(1171)	1792	0(1)	1(0)	1	38(73)	102(67)	150
BG/DPL	CA/UB	CA/UB/5	7	3	35	2	2	Master/1	714(0)	54(768)	768	85(96)	11(0)	96	0(0)	0(0)	0
BG/DPL	CA/UB	CA/UB/1	6	2	8	2	1	Master/1	564(0)	460(1024)	1024	25(10)	7(22)	32	35(0)	200(235)	240
BG/UF	ZYS/UB	ZYS/UB/5	8	1	12	2	3	Master/1	323(640)	317(0)	640	47(56)	17(8)	64	0(0)	0(0)	0
BG/UF	ZYS/UB	ZYS/UB/3	6	0	9	2	1	Master/1	1895(2048	153(0)	2048	0(0)	1(1)	1	0(0)	0(0)	0
BG/UF	ZYS/UB	ZYS/UB/4	7	1	10	2	1	Master/1	967(1920)	953(0)	1920	9(0)	7(16)	16	0(0)	0(0)	0
BG/UF	ZYS/UB	ZYS/UB/2	8	0	8	2	1	Master/1	952(322)	1096(1726	2048	0(0)	1(1)	1	0(0)	0(0)	0

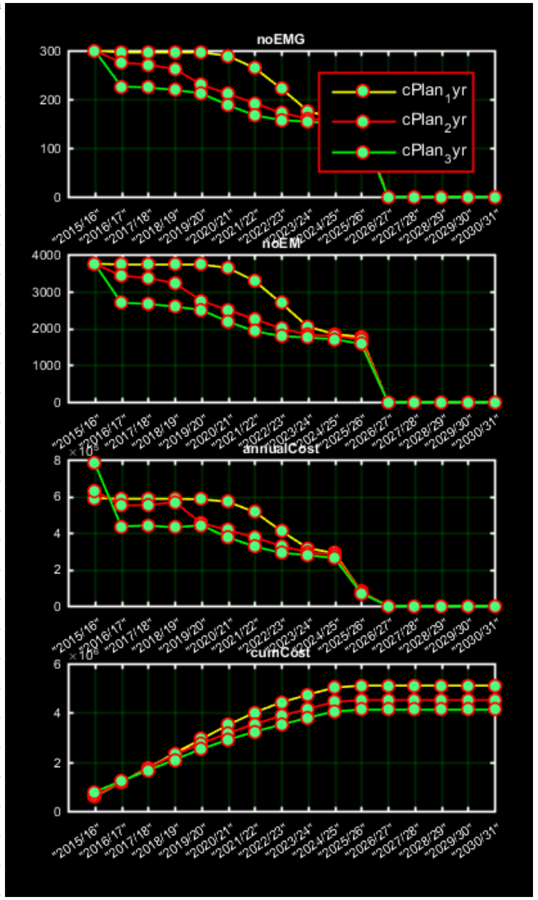

Current applications

Plan and build:

- AI steers investment
- AI optimises design
- AI assures outcomes
- AI mediate human goals

Service delivery and operations:

- Alarm management, anomaly detection
- Operational automation

Service creation and management:

- Watch this space!

Future applications

Creating value with AI in the Network – "enabling a zero touch network"

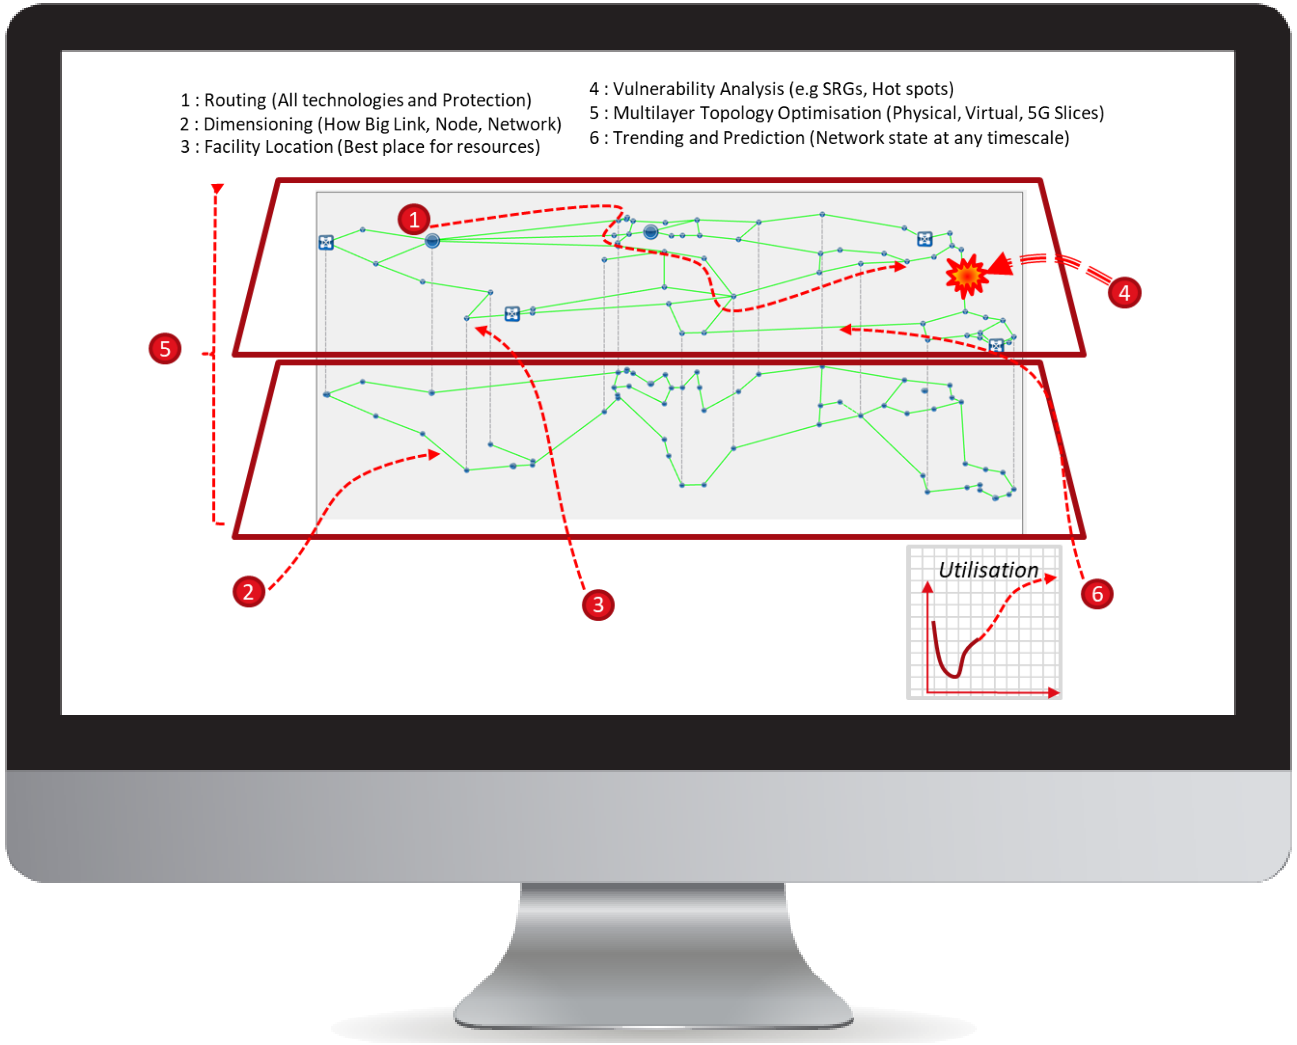

Case Study: BT – Supporting national growth through regional R&D partnerships

BT is funding AI research at 15 leading universities across the UK, and is the UK's largest telecoms and ICT investor in R&D. In addition, BT is leading a five year, £5m partnership with the Universities of Lancaster, Cambridge, Surrey and Bristol, as part of EPSRC's £78m Prosperity programme, creating an AI powered next generation data infrastructure for the UK. BT is expanding its global R&D centre and start-up cluster at Adastral Park in East Anglia.

Specifically, BT is committed to creating carrier scale and critical national infrastructure ready AI technology. This is an integral part of the evolving consortium of UK companies and institutions behind a proposed national Future Networks Research Centre with its hub at Adastral Park to drive AI into the global telecoms infrastructure.

New network technology creates new challenges for AI

vRAN Overview: base station architecture

- Load balancing – complex / adaptive choice due to transport constraints
- Compression trade-off / prediction
- Complex network design / heuristic

> - Base station consists of RRU (remote radio) and a BBU (baseband processor) which are connected via non-ideal transport
> - RRU and vBBU from different vendors
> - Open API between RRU and vBBU
> - RRU and vBBU reconfigurable (software on general purpose hardware)

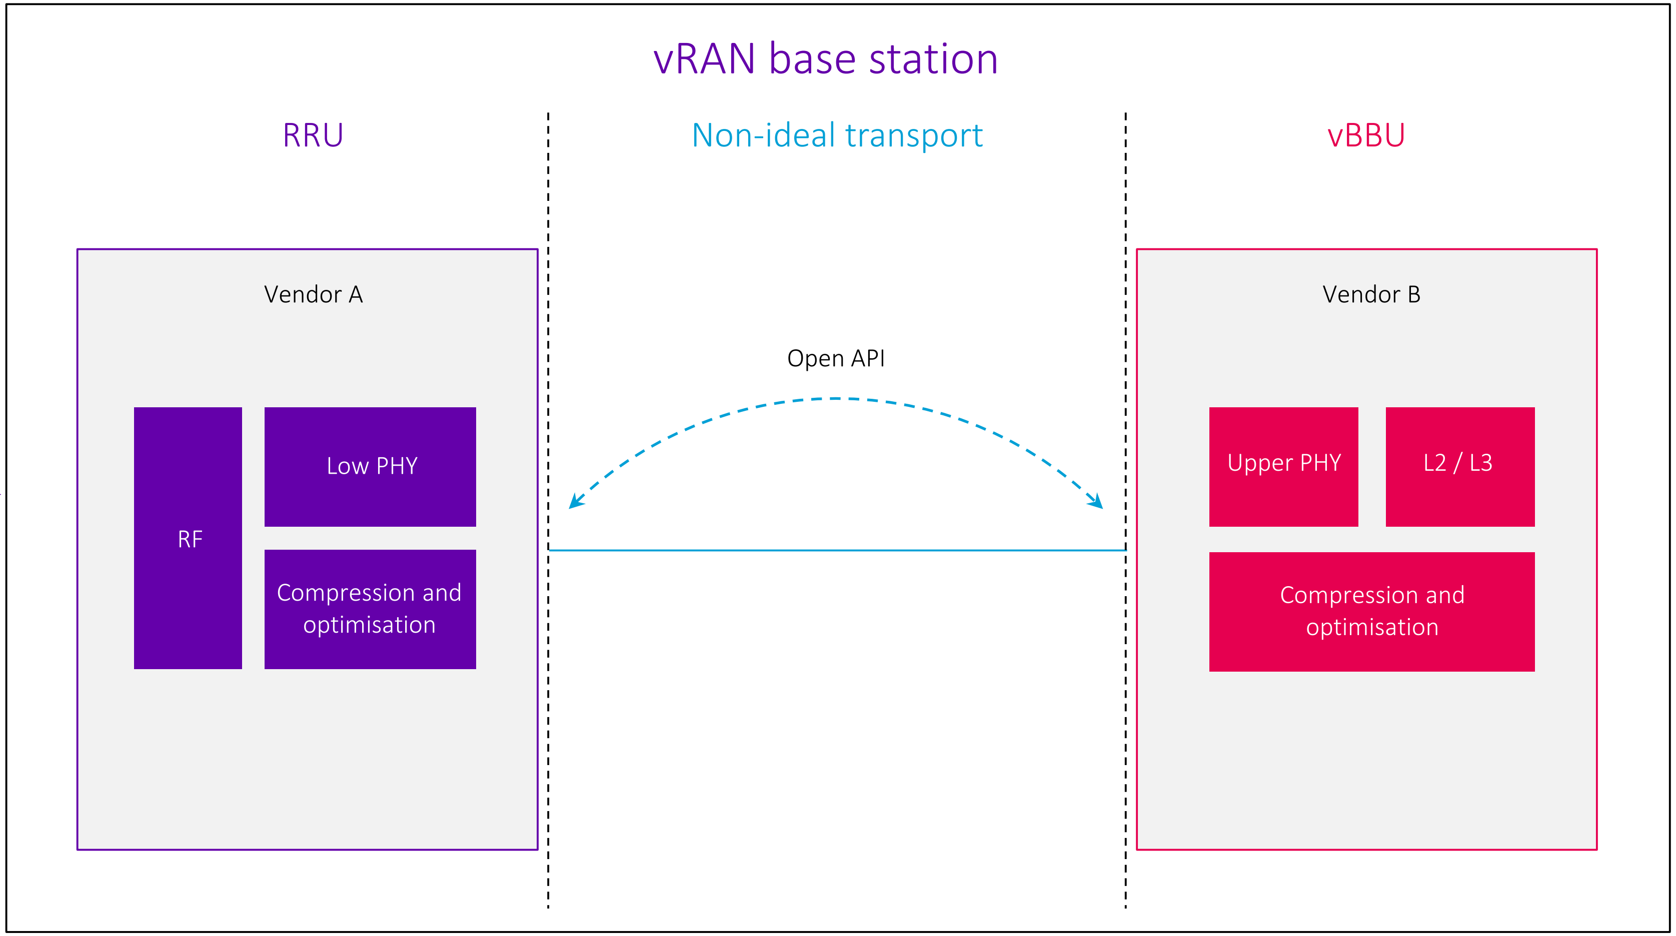

AI for Edge Compute

- How do we calculate service pricing to manage utilisation and offload?
- How do we understand performance of a component In the context of an end-to-end best effort service stack?
- How do we choose service configurations and manage the ongoing reconfiguration to get utilisation?

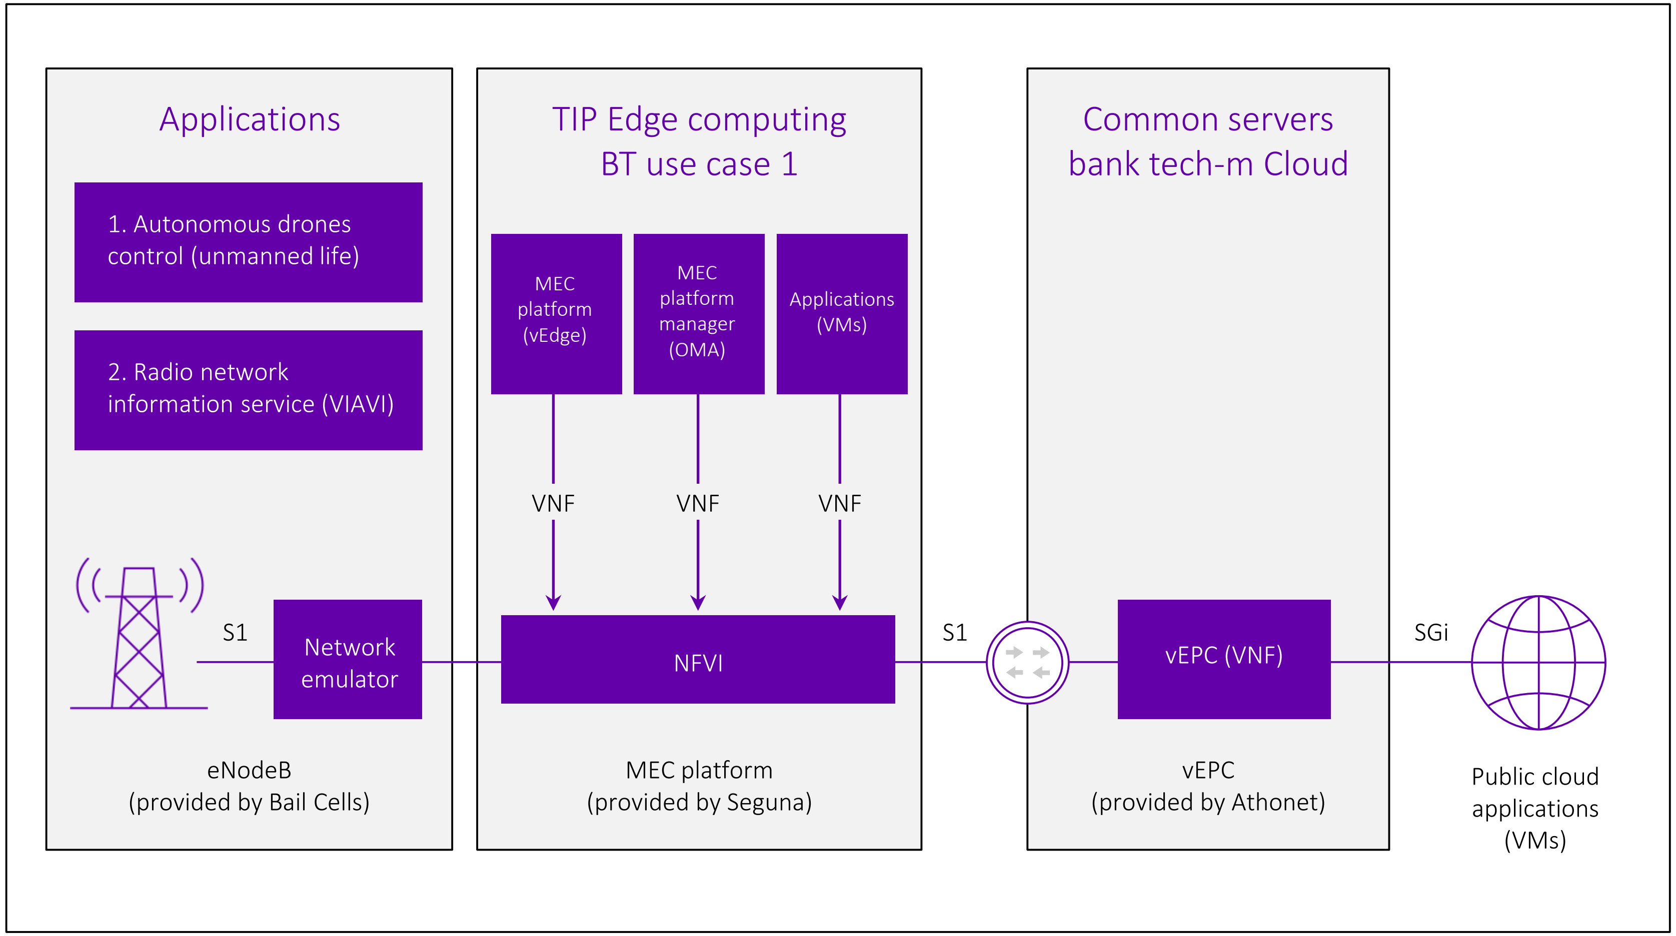

What's this all about again?

Next Generation Network platform – built around orchestration

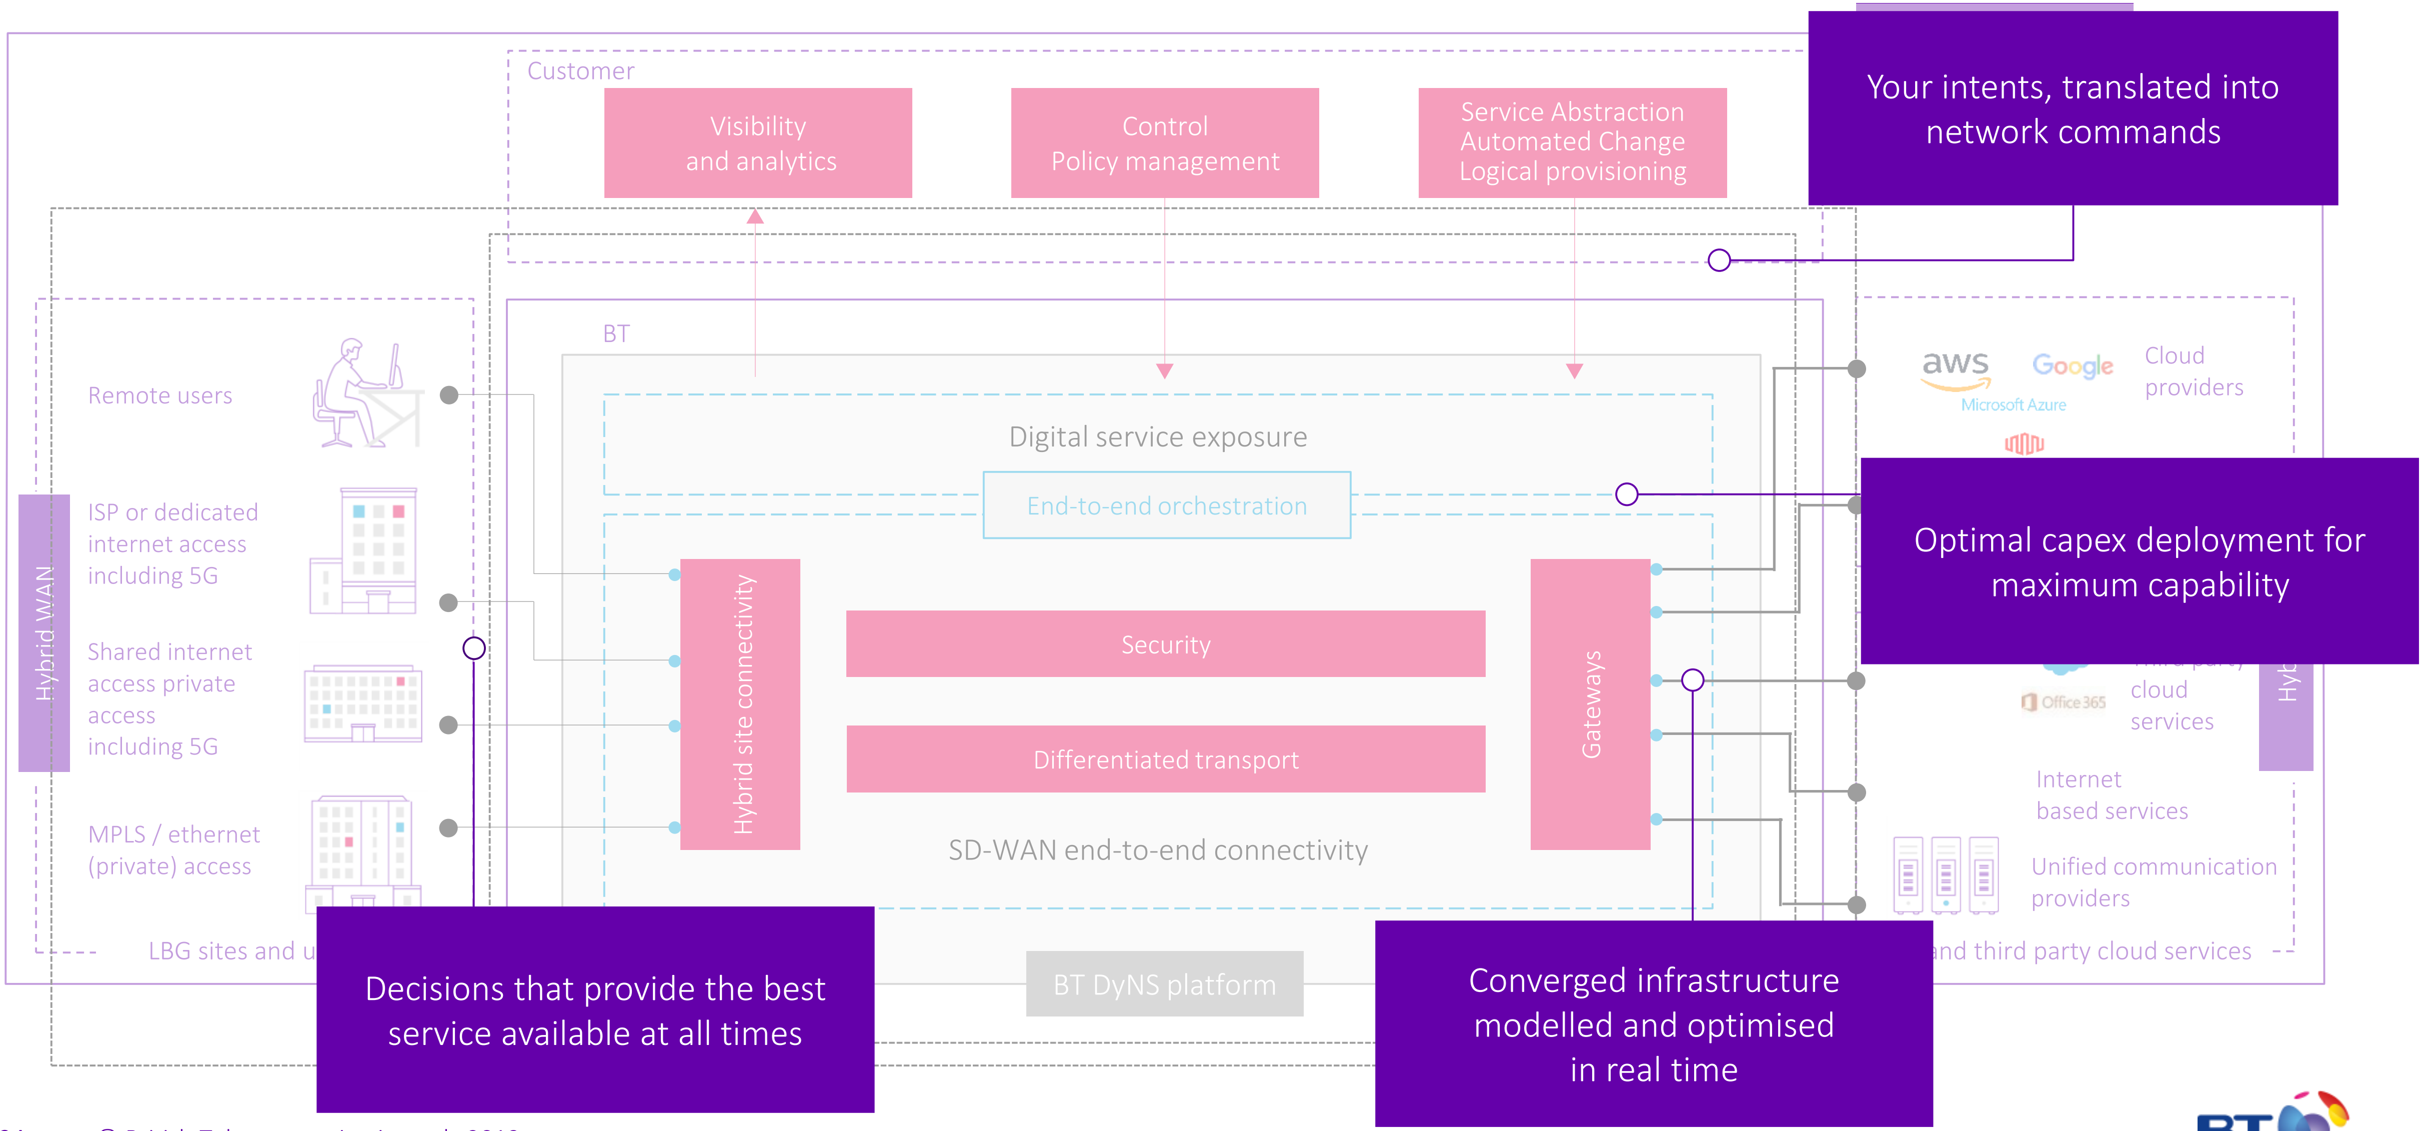

BT