Business Autonomics

David Rohlfing, Steve Cassidy,
Botond Virginas, Kjeld Jensen, Kashaf Khan

12th Dec 2017

Overview

Applying autonomics within a business environment

Intended to be generic

May be shaped from telecoms perspective and its inherent complexity

Already had presentations on

Self CHOP(E)

A specific example

What I won’t be talking about

Autonomic cars

Robots

Autonomic computing

Network infrastructure trends

Each successive technology increases the economic scale,
and broadens the market appeal
…and requires / drives a new model for the organisation

Increasing service breadth
Increasing customisation
Increasing number of elements
Increasing pace of change
…will drive complexity beyond current management models

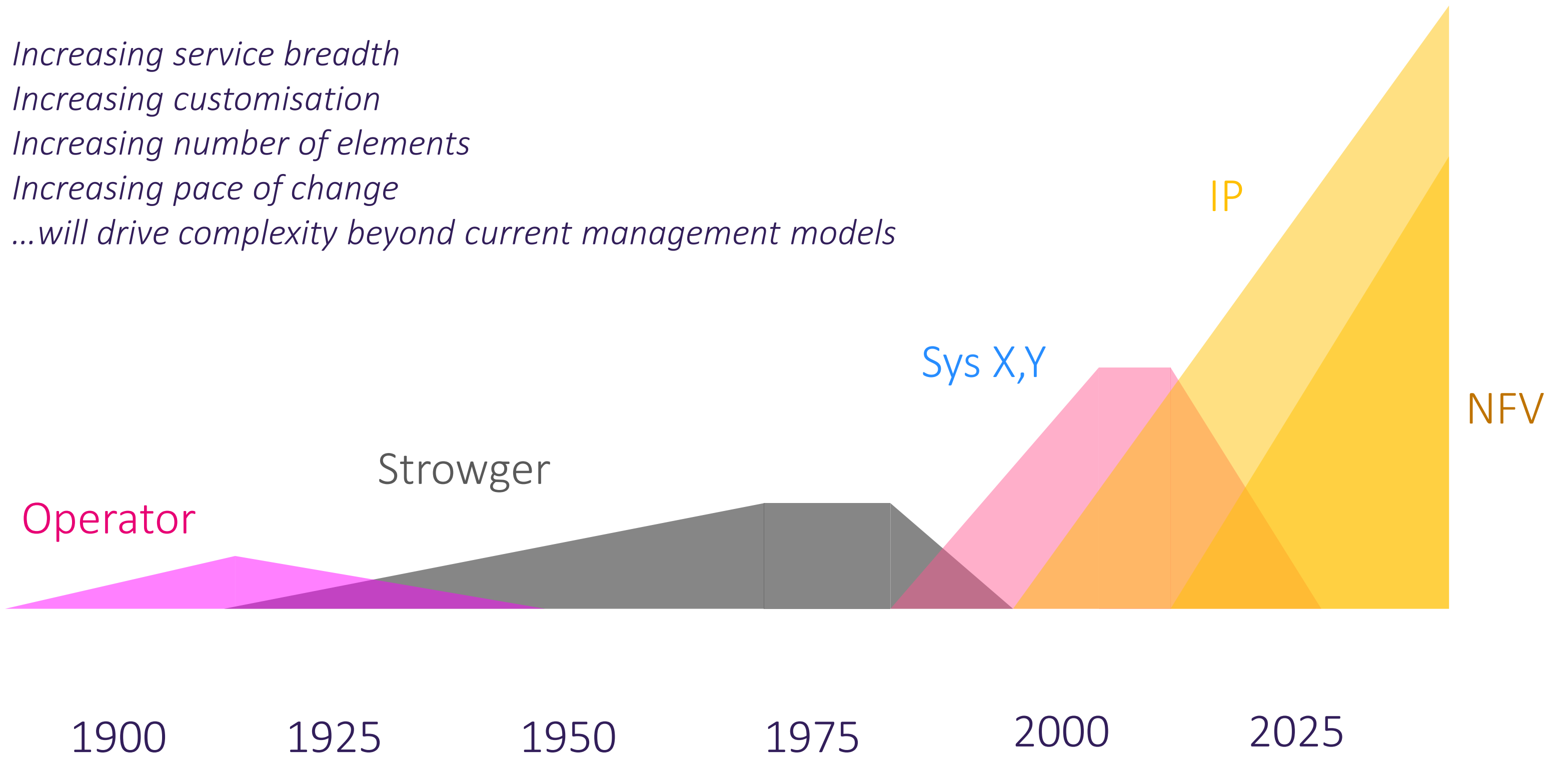

Coverage of "self-CHOP(E)" e.g. SON (mobile networks)

Configure	new base stations auto configure
Heal	minimises impact of failed nodes
Optimise	learning for base station control
Protect	proactive defence against unauthorised users and system attacks
Explain	Not yet present

How do we go about it ?

Collaborations

Universities

Industry

How do we go about it ?

Collaboration pt 1

Research questions

Where does it sit: the relationship with "AI"

The process of autonomic system design

Autonomic system loops

Special terms

"AI" constraints

Business Autonomous systems

Collaboration pt 2

Summary

Research Questions

What could go possibly wrong ??????

Recipe for assessment and "safe deployment areas"

Technical, Cultural (image problem?)

Suitability tool kit (possible? desirable?)

Definitions of: agility, robustness, learning risk… so can trade-off

Ability to explore balances / self-consistency, boundaries, test

What is enough? Limits of appropriateness.

Transferability of learning to other applications

Organisational scope and boundaries / remits / success criteria

Boundaries between autonomic systems… gaming

Human interface – what is normal? Ability to explain actions/status

System loads & penalties, speed/accuracy, thresholds/anomalies

Inclusion of context (outside the data) in decision processes

Exploration outside the "normal" scope of operation and responses (and how to put back into the system). Emergence.

Moving with the people – what's the evolution, skills, numbers, culture, structure

(systems adapt, what about the people adaption)

…

…

…

…

Key messages

System and business level description of the opportunity and cautions

Agnostic descriptions of what it is (&isn't)

What answers we are creating, to address what problems

Relationship with "AI" and other such – draw the full system jigsaw

What bits exist, in the business and in research

It's a *process*, which now changes what *design, test operations* are. Never settles.

Checklist of components/requirements – including of the estate, and organisation (Involve CIOs)

Start the people and organisation conversation (platforms and Group HR)

People-machine boundary (e.g. decision models vs process autonomy)

How do we go about it ?

Collaboration pt 1
Research questions
Where does it sit: the relationship with "AI"
The process of autonomic system design
Autonomic system loops
Special terms
"AI" constraints
Business Autonomous systems
Collaboration pt 2
Summary

Taxonomy

AI	➢Umbrella term – too general
Machine learning	➢Experts steer statistical models
Intelligent system	➢Information retrieval using context
AR	➢Retrieval based on location/activity
Decision support system	➢Data-driven storyboards retain context & precipitate cross-domain agreements
Expert (knowledge-based) system	➢Information regurgitation based on queries
Big data	➢Metaphor exploration
Data base	➢Fixed metaphor query support
Automation	➢Fixed recipe response driving actions
Autonomics	➢Closed loop – drives actions; double loop learning, CHOPE.

Machine-Learning is a small proportion of the *system* design and operational dynamics.

Levels of "AI"

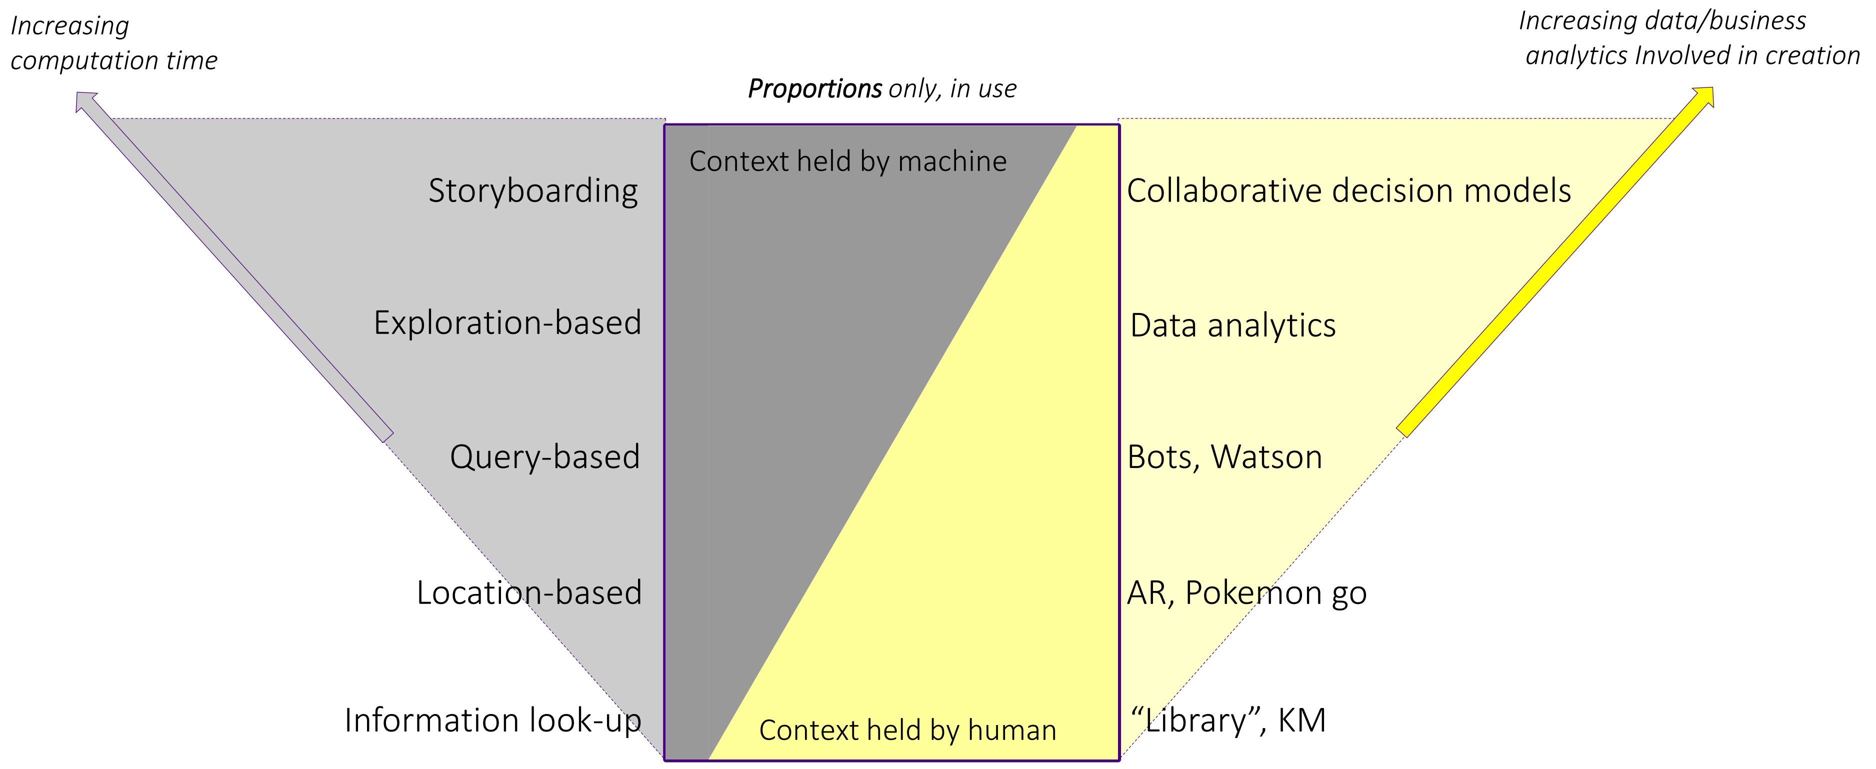

The process involved in an autonomic system is to continuously shuffle down the diagram from context-heavy, deep analytics and slower computation, (inc. forging agreements between stakeholders and domain experts), to faster more automated "reflex" decisions, using stripped-down images of the domain, with lower context.

How do we go about it ?

Collaboration pt 1
Research questions
Where does it sit: the relationship with “AI”
The process of autonomic system design
Autonomic system loops
Special terms
“AI” constraints
Business Autonomous systems
Collaboration pt 2
Summary

Autonomic system design is a *process*

An autonomic system ***cannot be "designed" in the traditional sense****. It is the result of a* ***continuing process*** *of iterative analysis with stakeholders, involving the discovery of real system limits and arriving at new agreements…*

The creation of the discriminator algorithm *requires* iteration with the domain expertise and the business as the "settings" which create the algorithm are discovered through iteration with the data, through which the possible balances between stakeholders are discovered and agreed.

The creation of the algorithm requires a further layer of business requirements to be made explicit; in addition to the direct product/service/cost requirements of the business case. These include required learning rate, risk, cost of learning.

The autonomic system, running on a (necessarily) simplified model of the real system, needs a further set of controls to ensure that safe operation is ensured. These are part of the design. (Beyond "AI").

All autonomic systems require a second learning loop, to ensure that changes in the real world, and further possible improvements to the existing system are capable of being supported.

The system will also need to support query interfaces etc. to ensure operational trust (this bit to come).

Next slides say a little more about the main 4 points here….

First stage

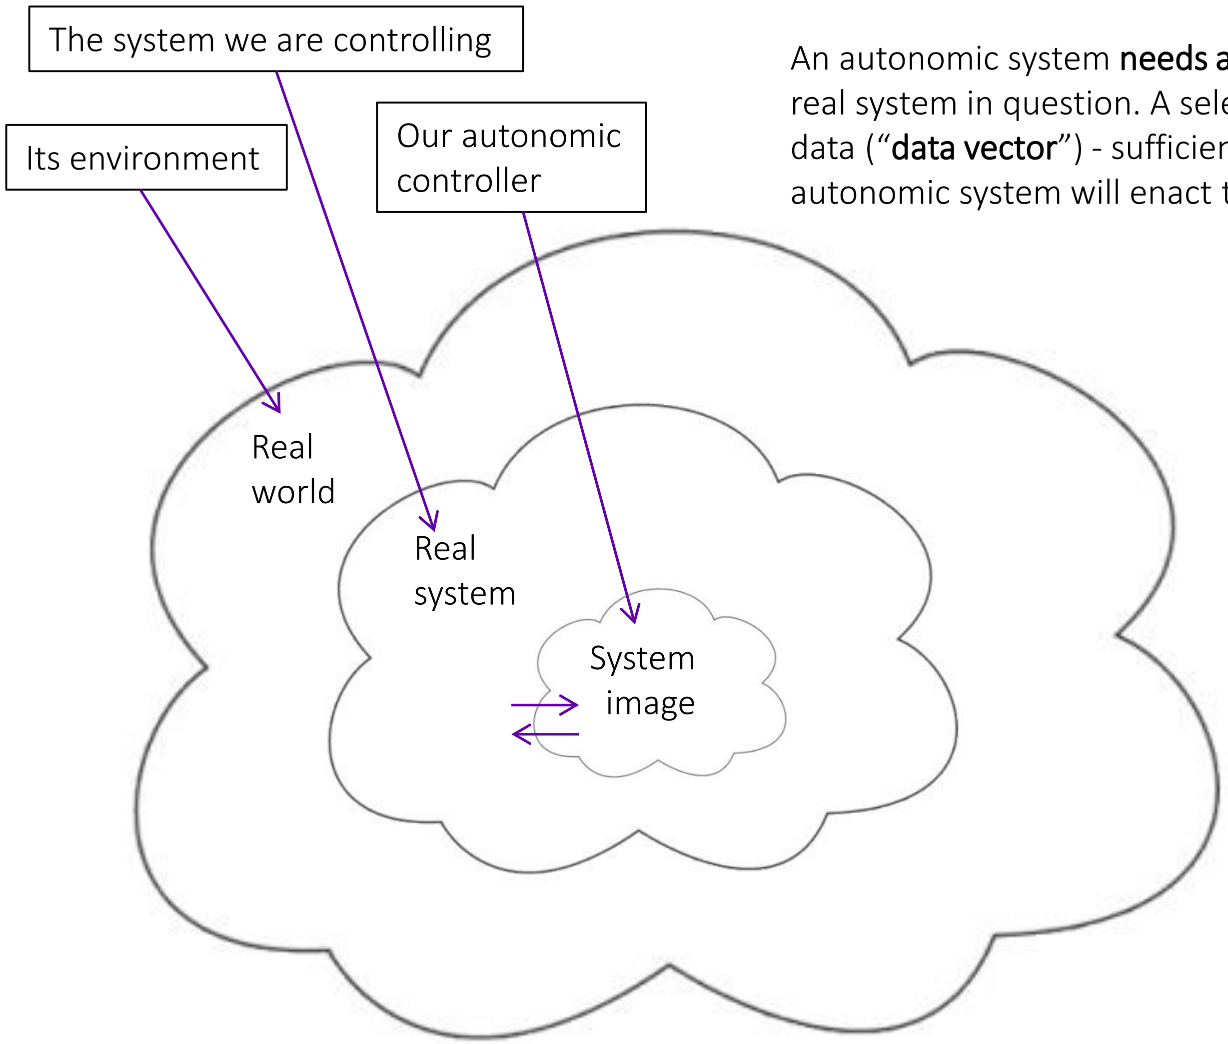

An autonomic system **needs a stripped-down image** of the real system in question. A selection of sufficiently representative data ("**data vector**") - sufficient to discriminate between actions the autonomic system will enact to support desired outcomes.

Process:
Analysis to determine **how accurately** the real system can be **represented** by a system image.

Does the available data describe the system dynamic well enough?
Hence determine the "image vector"- the (data) projection of the real system into the system model (algorithm).

Data analysis, stats, domain expertise.
Exploratory, iterative.

The creation of the **algorithm *requires* iteration** with the business domain as the stakeholders discover in the data the extent to which their **requirements need to be traded** to enable a working system model. (e.g. cost/service levels)

Second stage

Building a working system image *forces* **explicit agreement on learning rates, accuracy / risk, cost of learning**. These are parameters necessary in the creation of the algorithm.

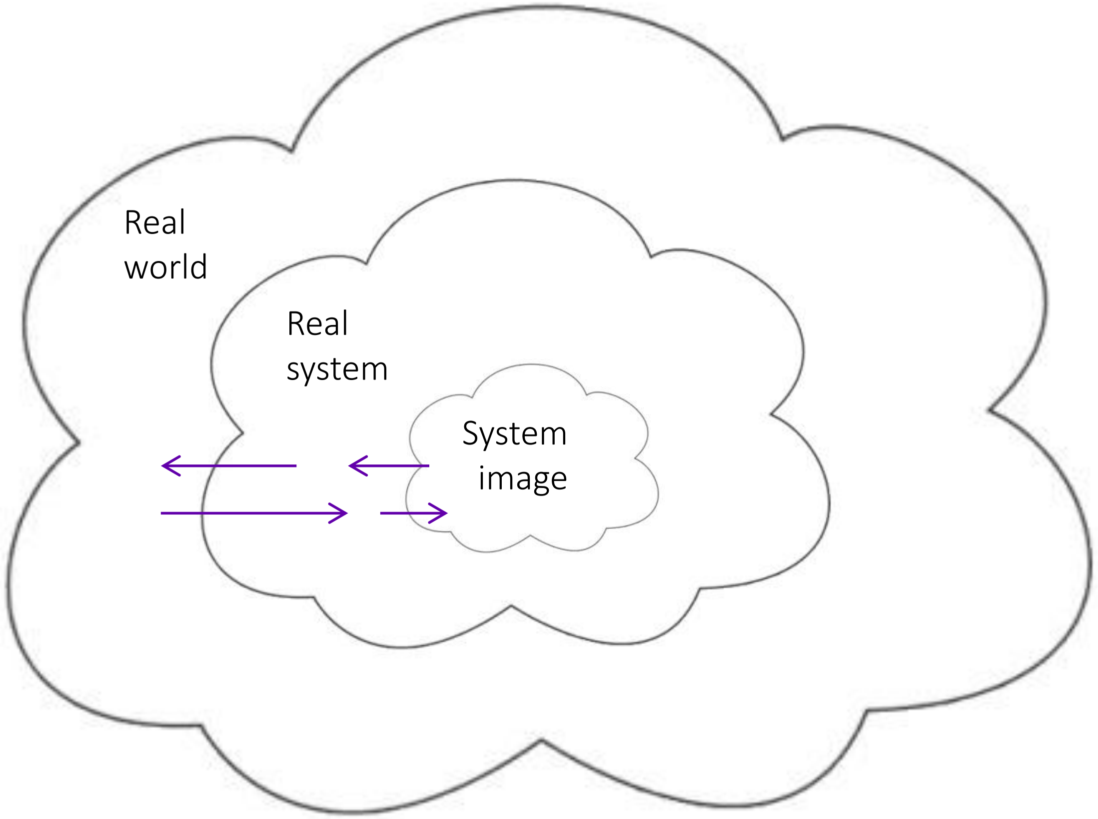

Process:
Analysis to determine the extent to which the real system is likely to be perturbed by the real world – **dynamics of the system and changes in business requirements**.

Design of learning strategies, including the ability to amend the decisions taken by the algorithm's discriminator output.

Testing of scenarios for change rates, including stress cases such as data errors, demand surges, outages.

Data analysis, stats, domain expertise. Modelling. Exploratory, iterative.

The creation of the algorithm *requires* **iteration** with the **business domain** to determine these "meta-requirements" through iteration with the data.

Third stage

To **protect the real system** from any "naïve" decisions from the "system image", various decision gates need to be interposed between the discriminator and the real system.

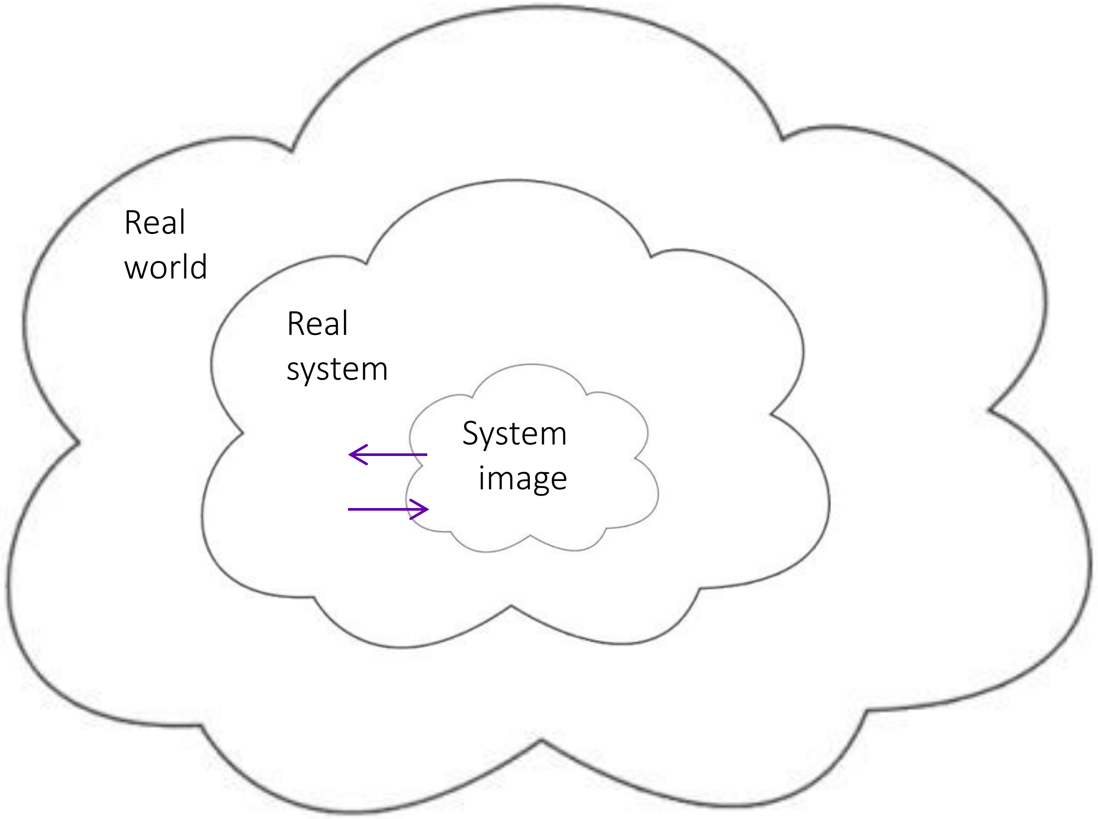

Process:
Determine the limits of decisions that the real system can absorb successfully (e.g. number of engineers sent)
Determine the actions to be taken when incoming data or learning rates change beyond thresholds.

Testing of scenarios for dynamic behaviours.
Assess "shadow system" performance.

Data analysis, stats, domain expertise.
Modelling. Exploratory, iterative.

The creation of the algorithm *requires* **iteration** with the **domain experts** (inc business) to determine the cost/benefit/feasibility **balances in (real) system costs/risks and investment** in learning and agility.

Fourth stage

Exploration is required to maintain an **understanding of the wider environment** in the real world as it continues to evolve. It may be necessary to change the image vector (e.g. further categories of data or discrimination; changes in expected dynamic etc.)

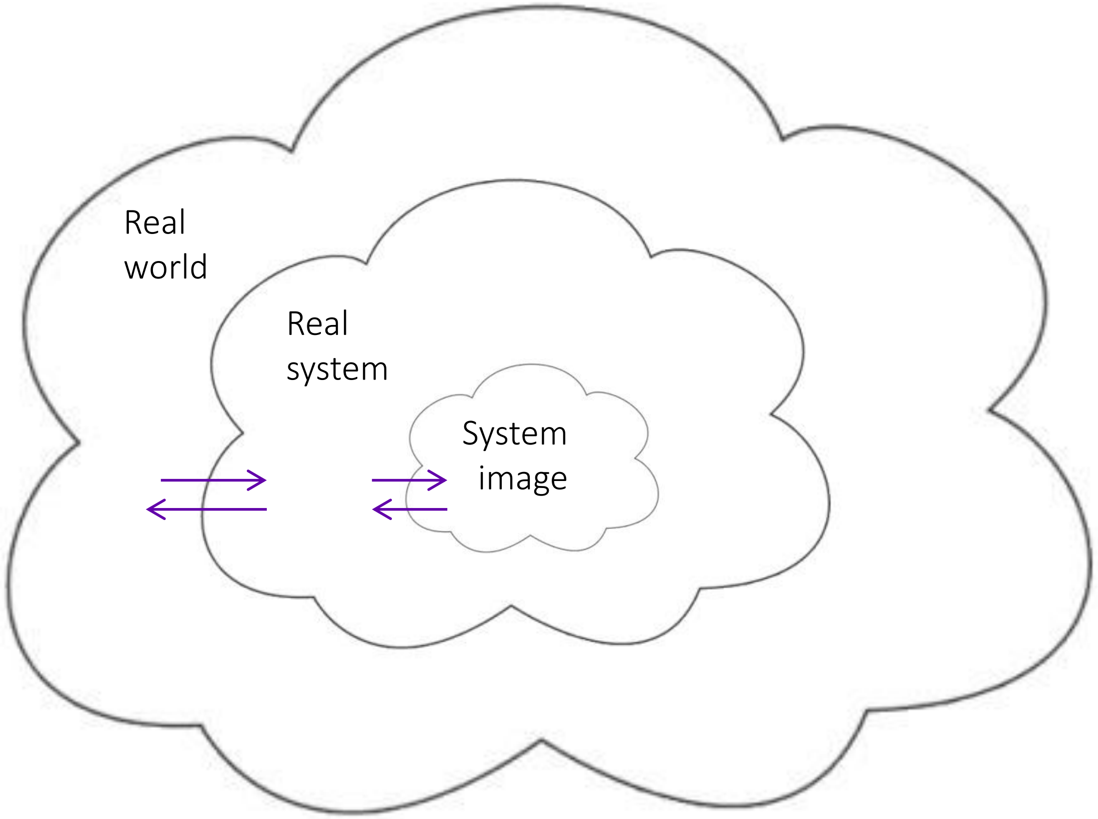

Process:
Revisiting stages 1-3 as a "2nd loop" activity.

How do we go about it ?

Collaboration pt 1
Research questions
Where does it sit: the relationship with “AI”
The process of autonomic system design
Autonomic system loops
Special terms
“AI” constraints
Business Autonomous systems
Collaboration pt 2
Summary

Autonomic system – loop 1

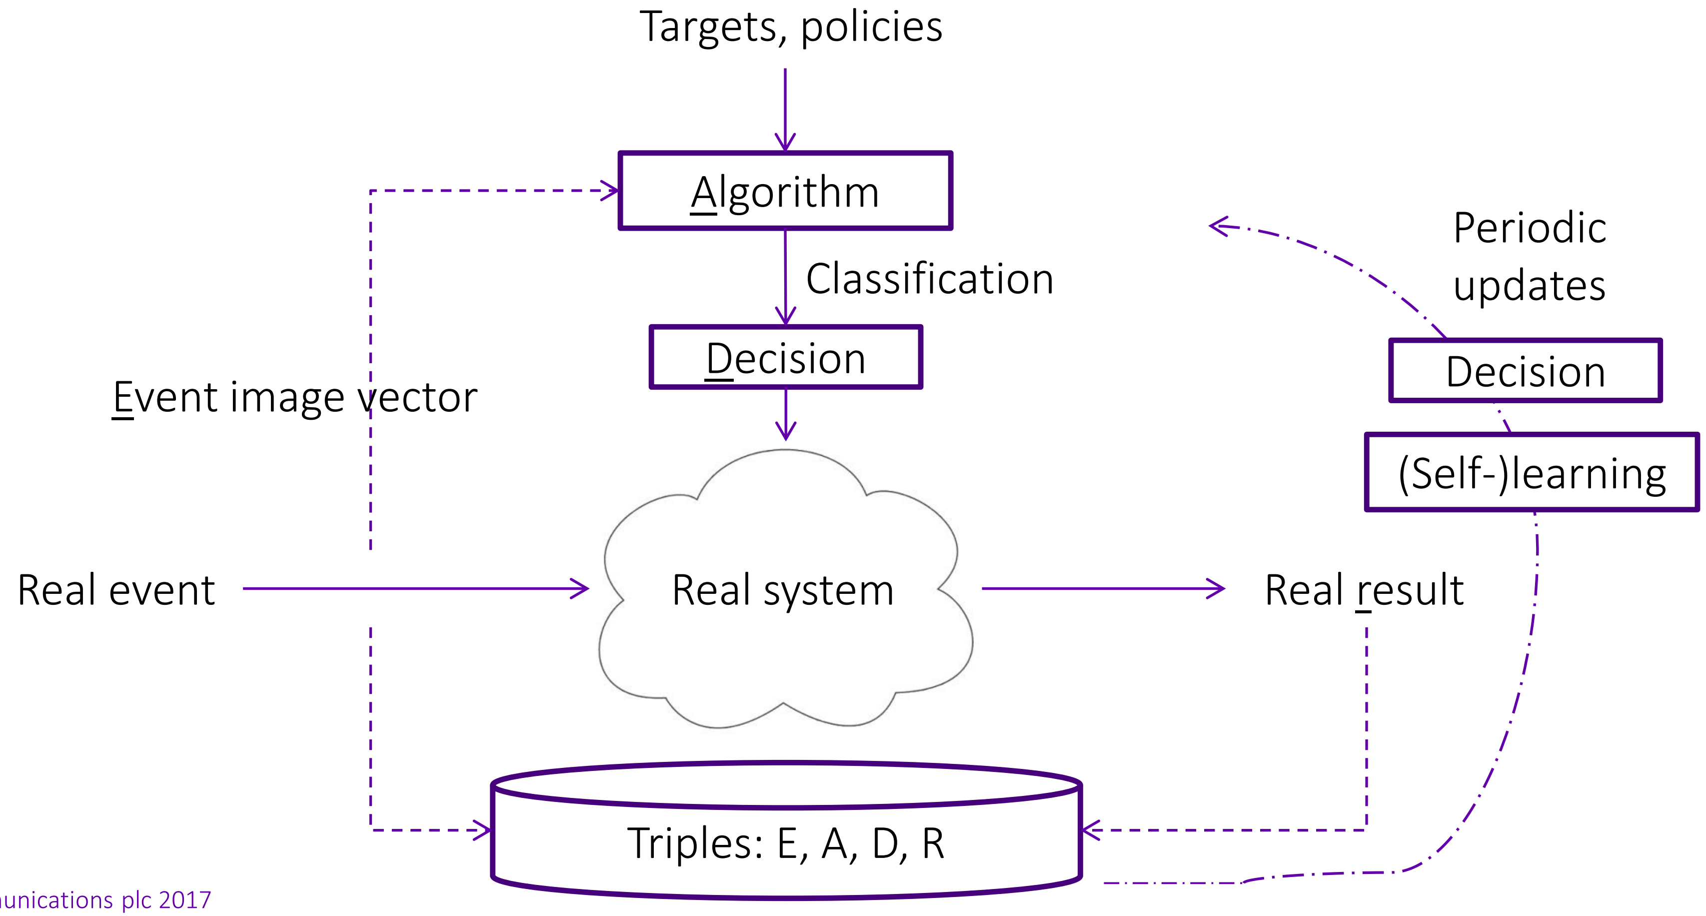

Autonomic system – loops 1 and 2

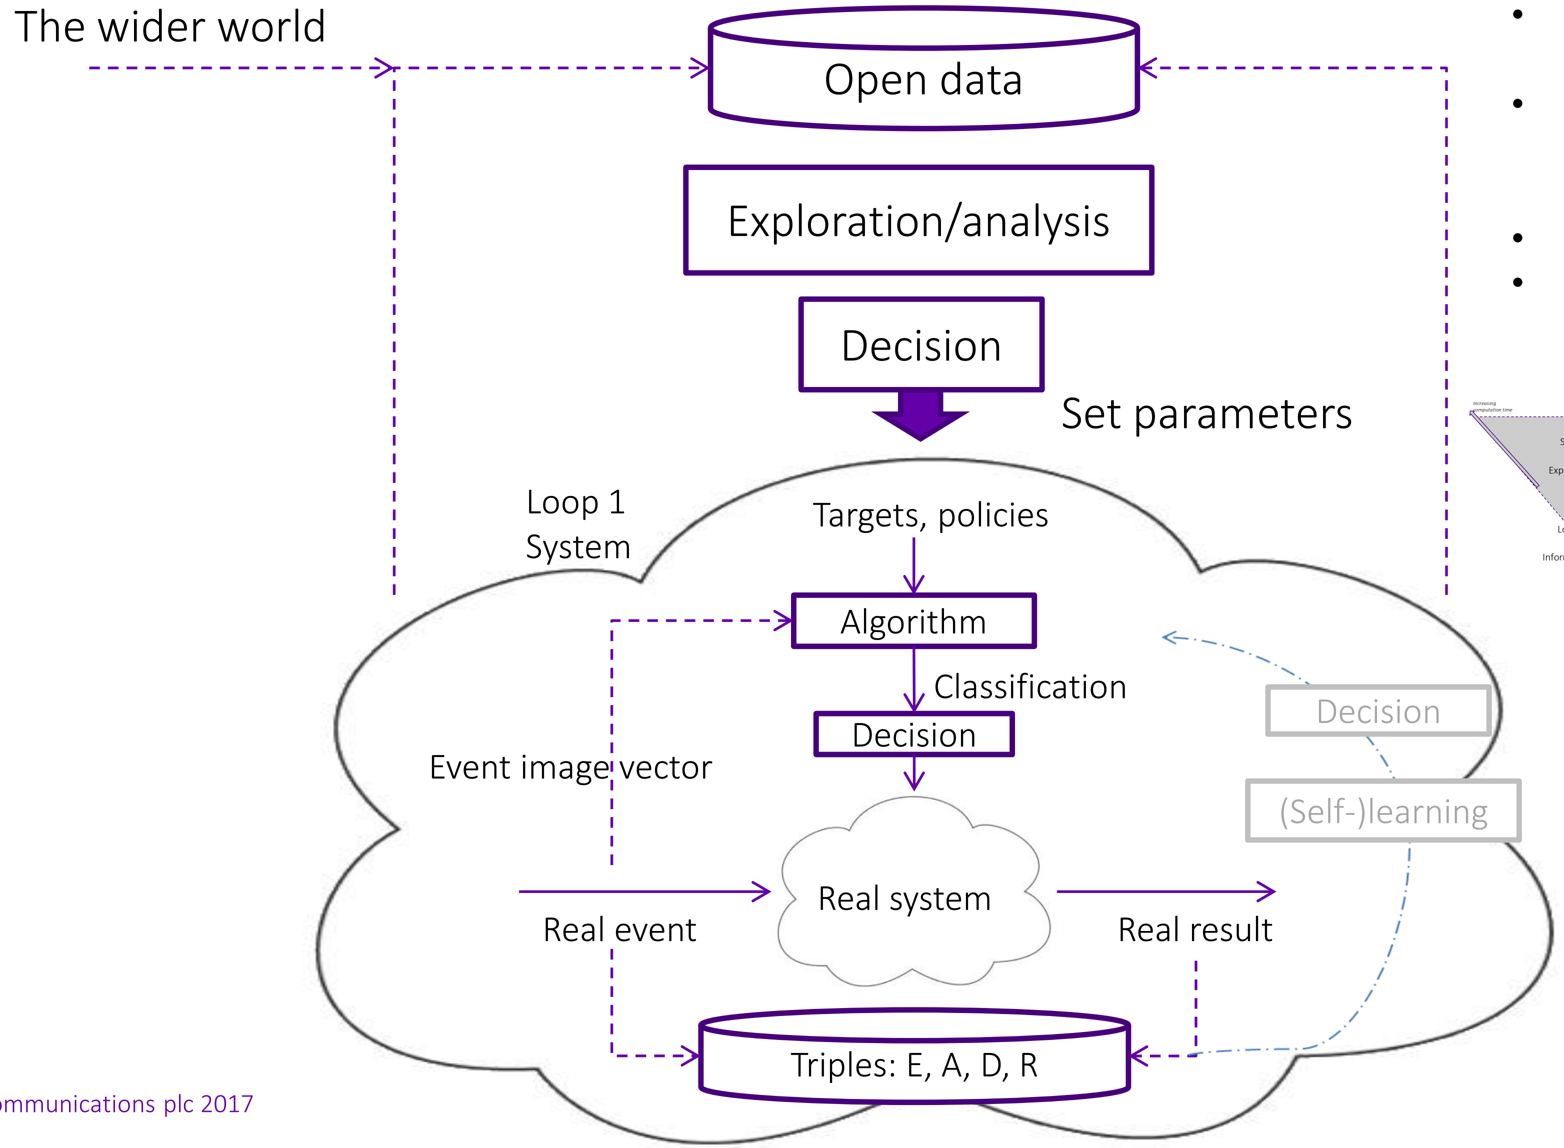

- Multiple stakeholders.
- Several possible objectives.
- Agree trade-offs.
- Proactive as well as reactive.

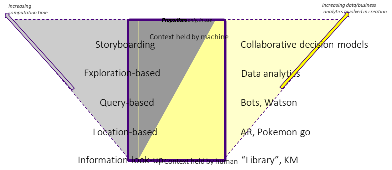

Kjeld/Botond groupings

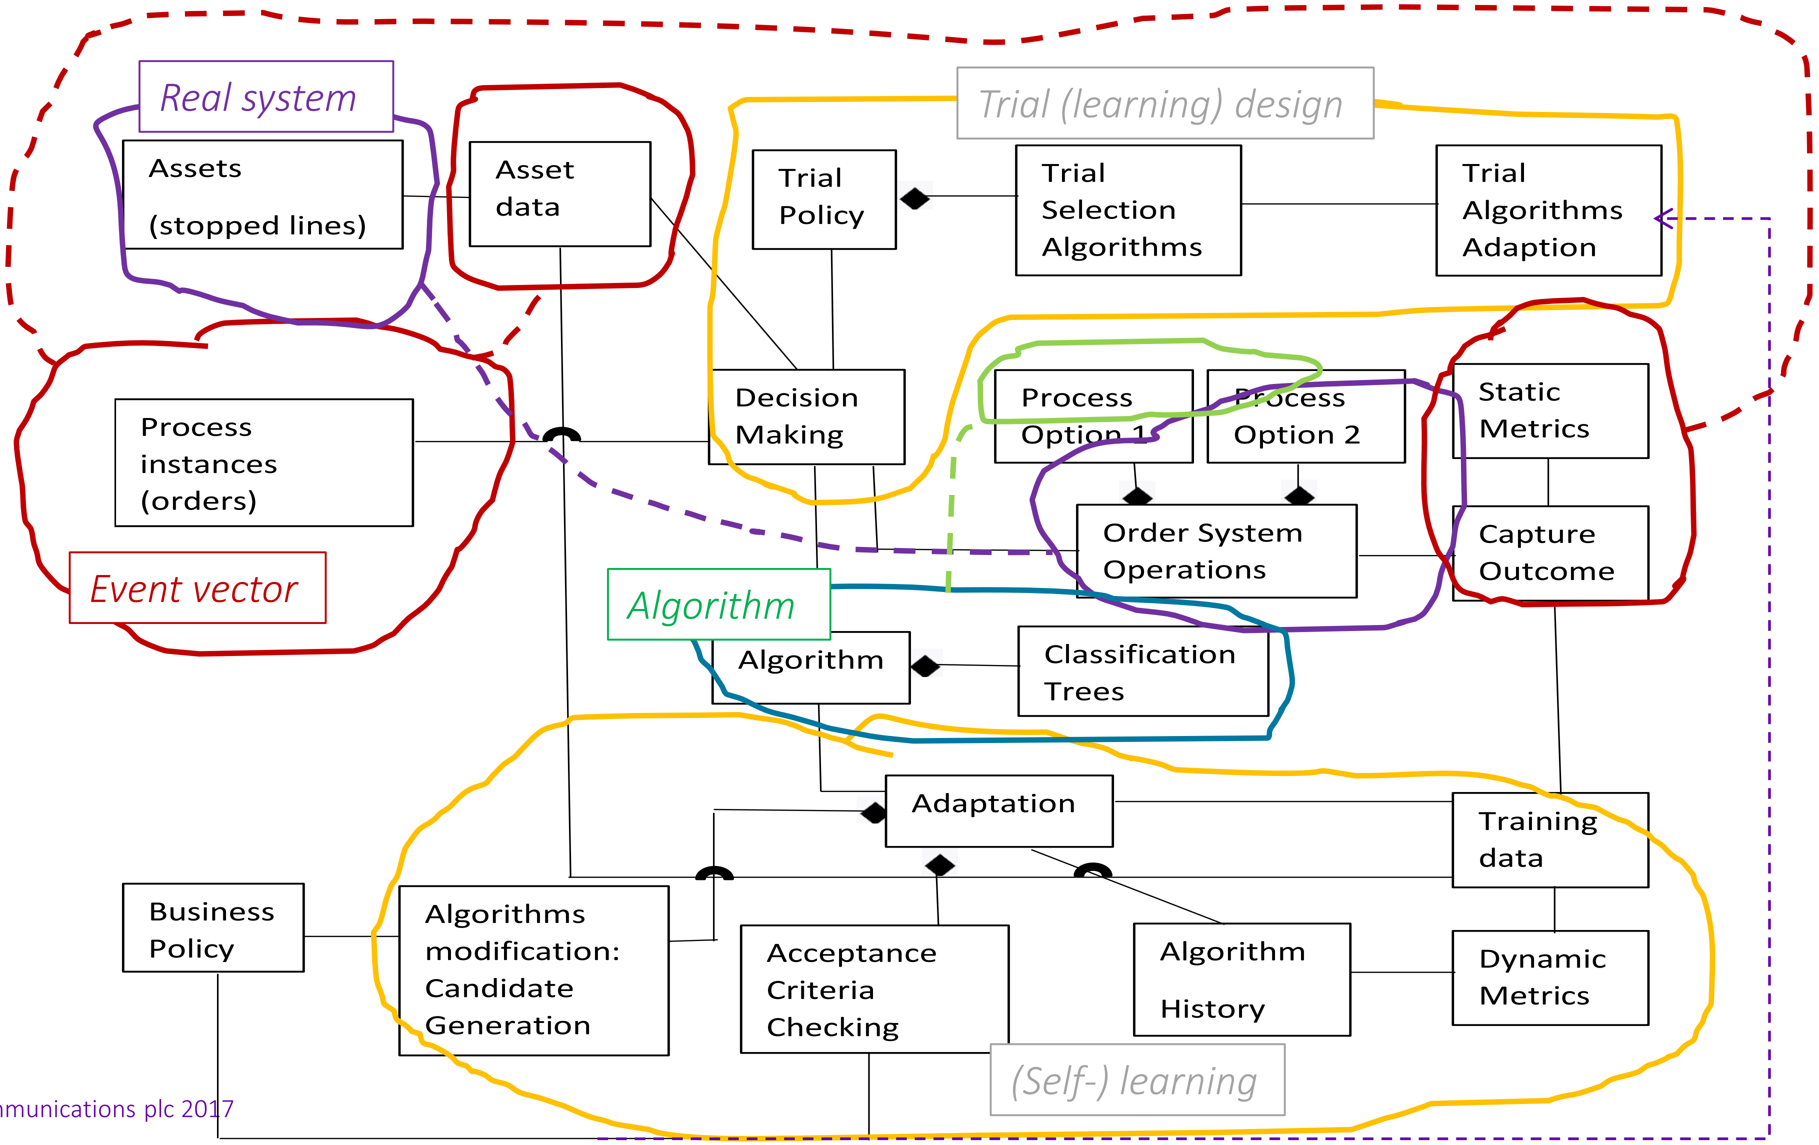

How do we go about it ?

Collaboration pt 1
Research questions
Where does it sit: the relationship with “AI”
The process of autonomic system design
Autonomic system loops
Special terms
“AI” constraints
Business Autonomous systems
Collaboration pt 2
Summary

Dictionary

Component	Definition	Subcomponents	Definition
Real world	The business environment: external markets, competition, and internal pan-business drivers	Stakeholders Portfolio Economy Service	
Real System	The system to be controlled: a unit of value generation eg a service process	Assets	The elements that provide data, or are the subject of *control actions*
		Control actions	The set of actions which *the real system* is able to enact
System image (controller)	The model of the system used to control the real system	Event image vector	The data set chosen as a sufficiently accurate representation of the *real system* (events, *control actions*, consequences)
		Control algorithm (discriminator)	The software that discriminates between potential *control actions* to be taken on the *real system*
		Decision (control)	The component that makes the choice of *control action*, based on the *algorithm's* discrimination, modified by *real system* protection or learning requirements
Learning system	The (off-line) system that updates the control algorithm	(self-)Learning	Building of new or modified *control algorithms*
		Decision (learning)	Enacts the choice of whether to implement a modification to the *control algorithm*
		Exploration	The 2nd loop learning which creates the 1st loop and maintains its relevance.

How do we go about it ?

Collaboration pt 1
Research questions
Where does it sit: the relationship with “AI”
The process of autonomic system design
Autonomic system loops
Special terms
“AI” constraints
Business Autonomous systems
Collaboration pt 2
Summary

Constraints in the application of “AI”

10^12 bits

>>10^36 bits

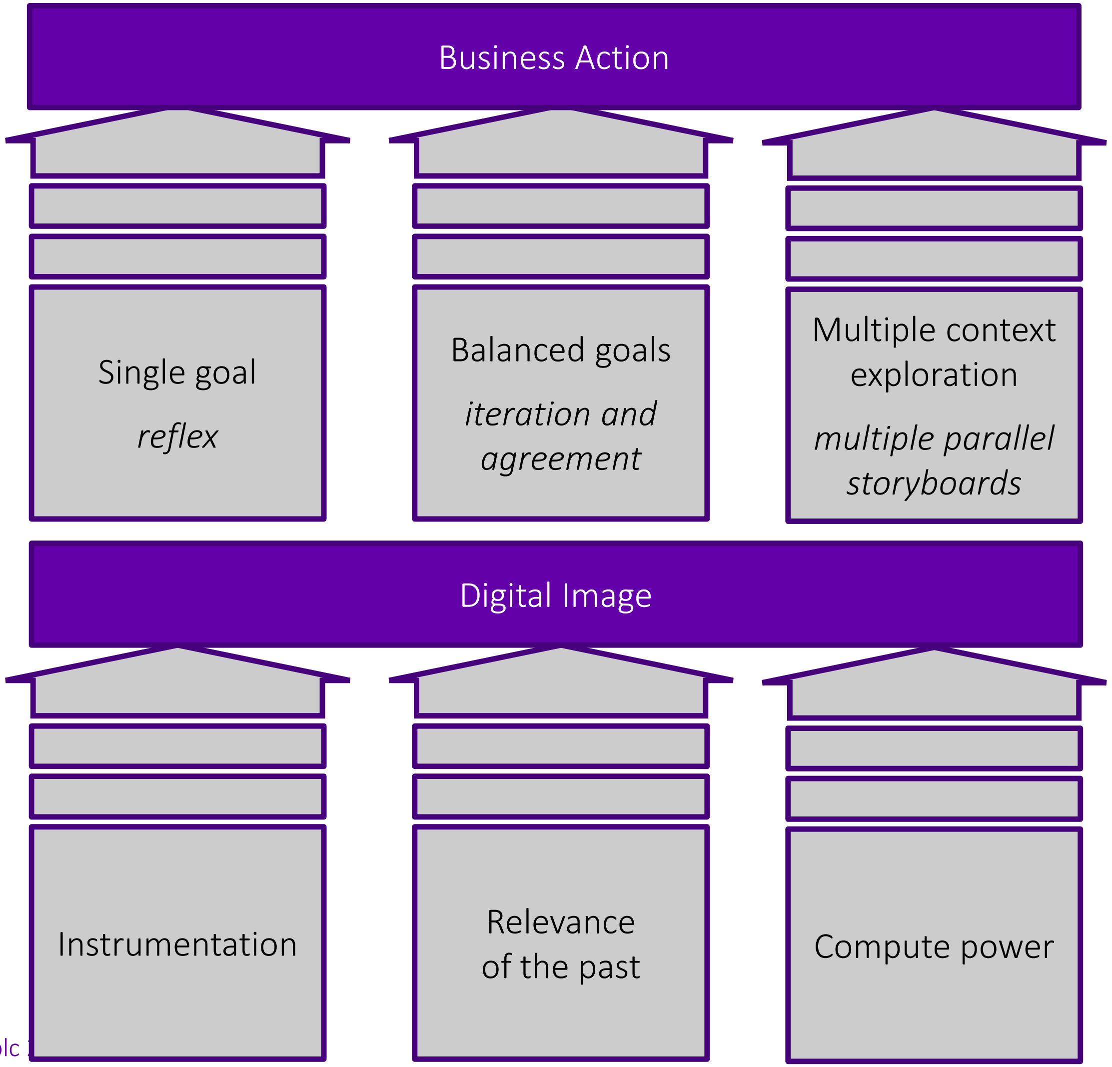

Search of the image space requires very different involvement of human insight. (extent of paradigm-based data collapse) Different human processes.

Fidelity of the image limited by the min of these. Investment is in domain-focussed analysis. (time, learning-rate limit)

Constraints in the application of “AI”- creation of the Digital Image

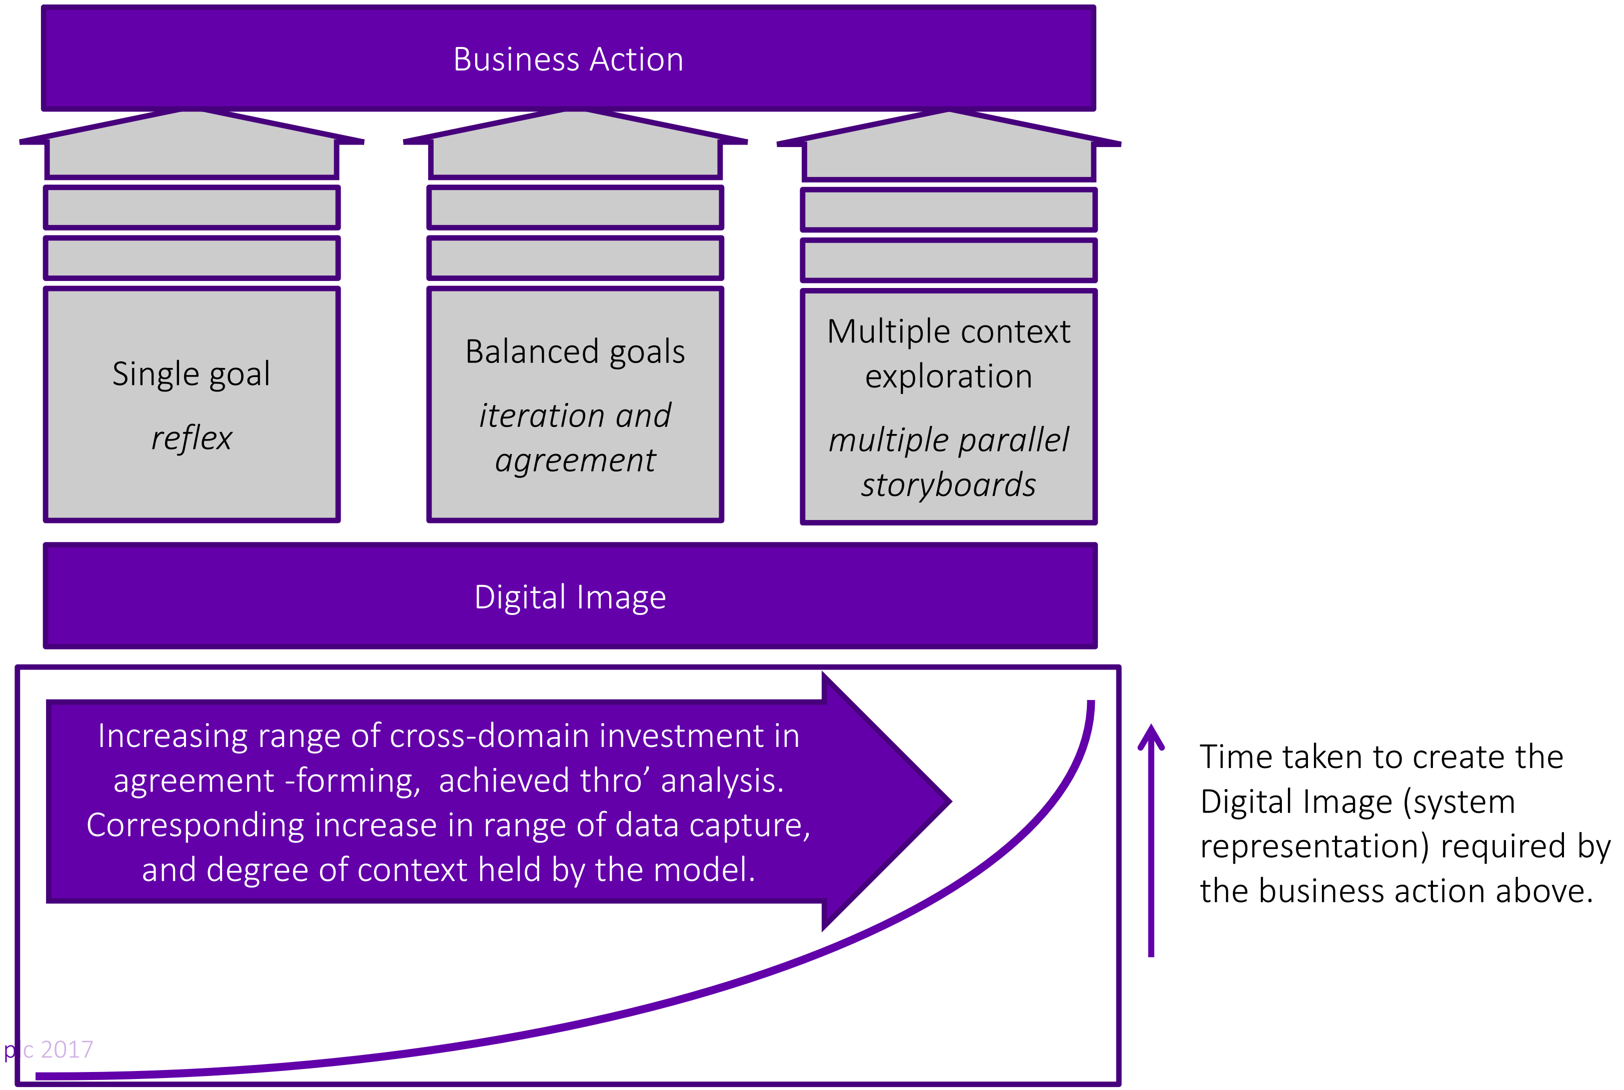

Constraints in the application of "AI" – searching the space of solutions

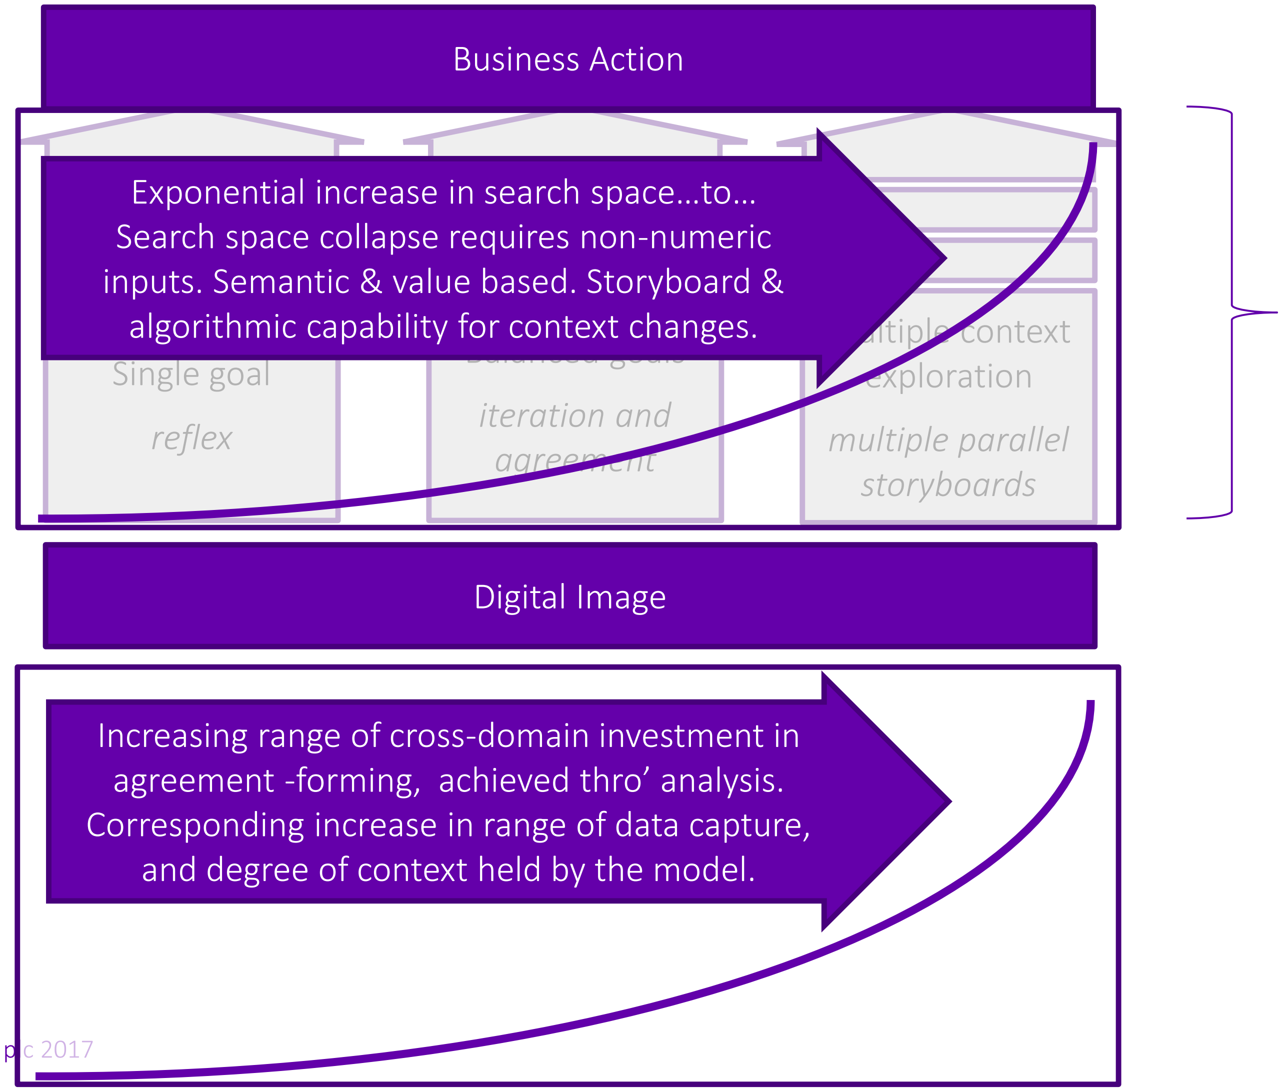

Search of the image space requires very different involvement of human insight. (extent of paradigm-based data collapse) Different human processes.

How do we go about it ?

Collaboration pt 1
Research questions
Where does it sit: the relationship with “AI”
The process of autonomic system design
Autonomic system loops
Special terms
“AI” constraints
Autonomous systems and the organisation
Collaboration pt 2
Summary

Autonomous systems – linkages with the organisation

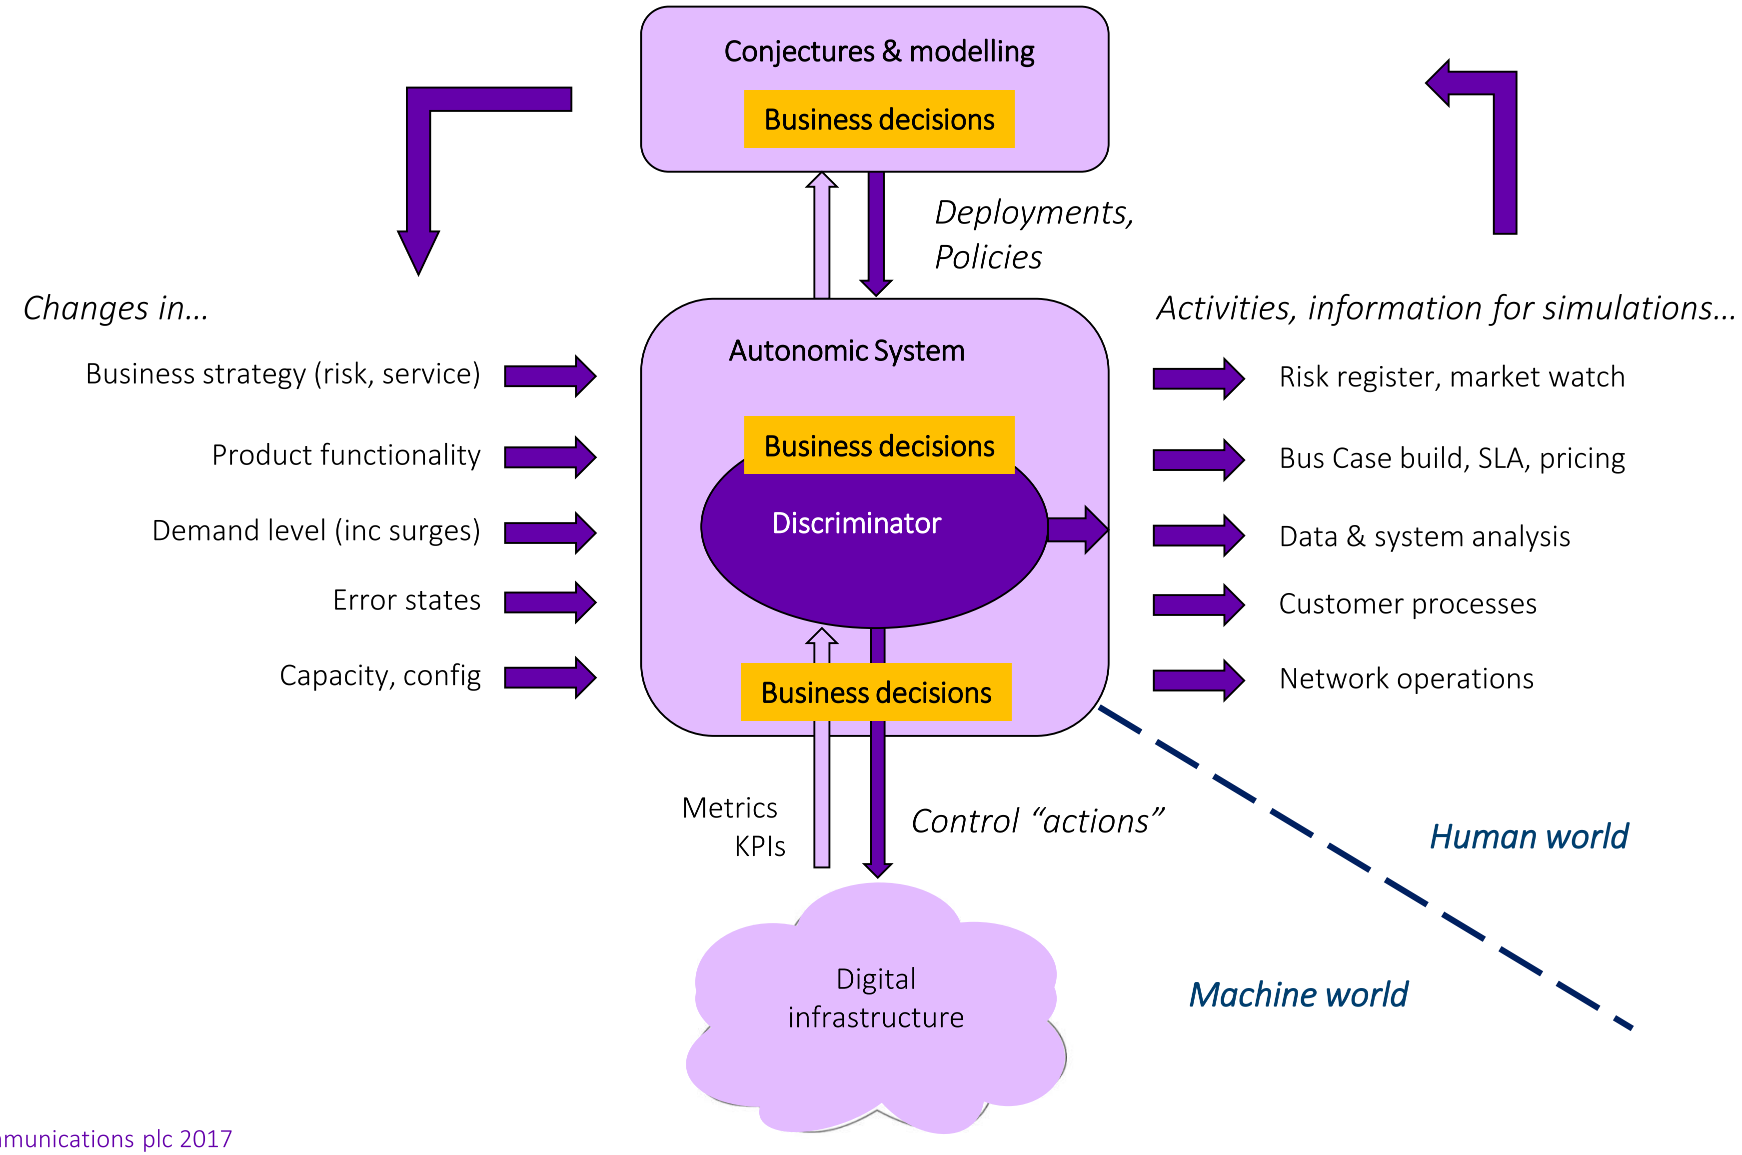

How do we go about it ?

Collaboration pt 1
Research questions
Where does it sit: the relationship with “AI”
The process of autonomic system design
Autonomic system loops
Special terms
“AI” constraints
Business Autonomous systems
Collaboration pt 2
Summary

Prosperity vision

Our team of internationally-renowned scientists and engineers working across top universities and BT will deliver the next generation converged digital infrastructure for the UK, by combining programmable and virtualised network functions with cutting edge data science and autonomous decision-making.

The resulting infrastructure will be reliable, resilient to disruptions, and able to cope with increasing and unpredicted demands on its capacity and types of services, to ensure that the UK's digital infrastructure continues to be world-leading.

Prosperity

This vision will be realised through the pursuit of three interdependent objectives:

1) Developing a completely new architecture for digital infrastructures, composed of highly dynamic network functions based on a micro-NFV approach that are collectively able to adapt to the real-time requirements of future digital services.

2) Creating a new autonomic framework for digital infrastructure to equip the nodes of the infrastructure network with the ability to understand their state, detect and diagnose disruptions to service, and take autonomous actions.

3) Implementing approaches for the successful integration of these technologies within the business functions with an aim to improve service assurance and organisational value.

Prosperity research challenges

Research challenge	Question	Outcomes
RC1: Agile converged infrastructure systems architecture [Hutchison, Willis]	How to build an agile digital infrastructure that is amenable for autonomous operation?	• Micro NFV, SDN • Generation of performance data
RC2: Future network operations and services [Tafazolli, Crane]	How to ensure service reliability of the agile autonomic digital infrastructure?	• Guaranteed QoS/QoE • Automated service ecosystem
RC3: Autonomic knowledge framework [McFarlane, Cassidy]	How to enable autonomous operational ability for the digital infrastructure?	• Multi-agent systems • Intelligent data
RC4: Autonomous diagnostics & response [Eckley, Cassidy]	How can the digital infrastructure respond to disruptions autonomously?	• Automated disruption detection • Automatic response learning
RC5: Future organisational dynamics [Parlikad, Cassidy]	How to exploit the autonomic agile capabilities of the digital infrastructure?	• Autonomous decisions • Organisational learning

University research challenge leaders

- RC1 Hutchison, Lancaster University
- RC2 Tafazolli, Surrey University
- RC3 McFarlane, Cambridge University
- RC4 Eckley, Lancaster University
- RC5 Parlikad, Cambridge University

CDI Components

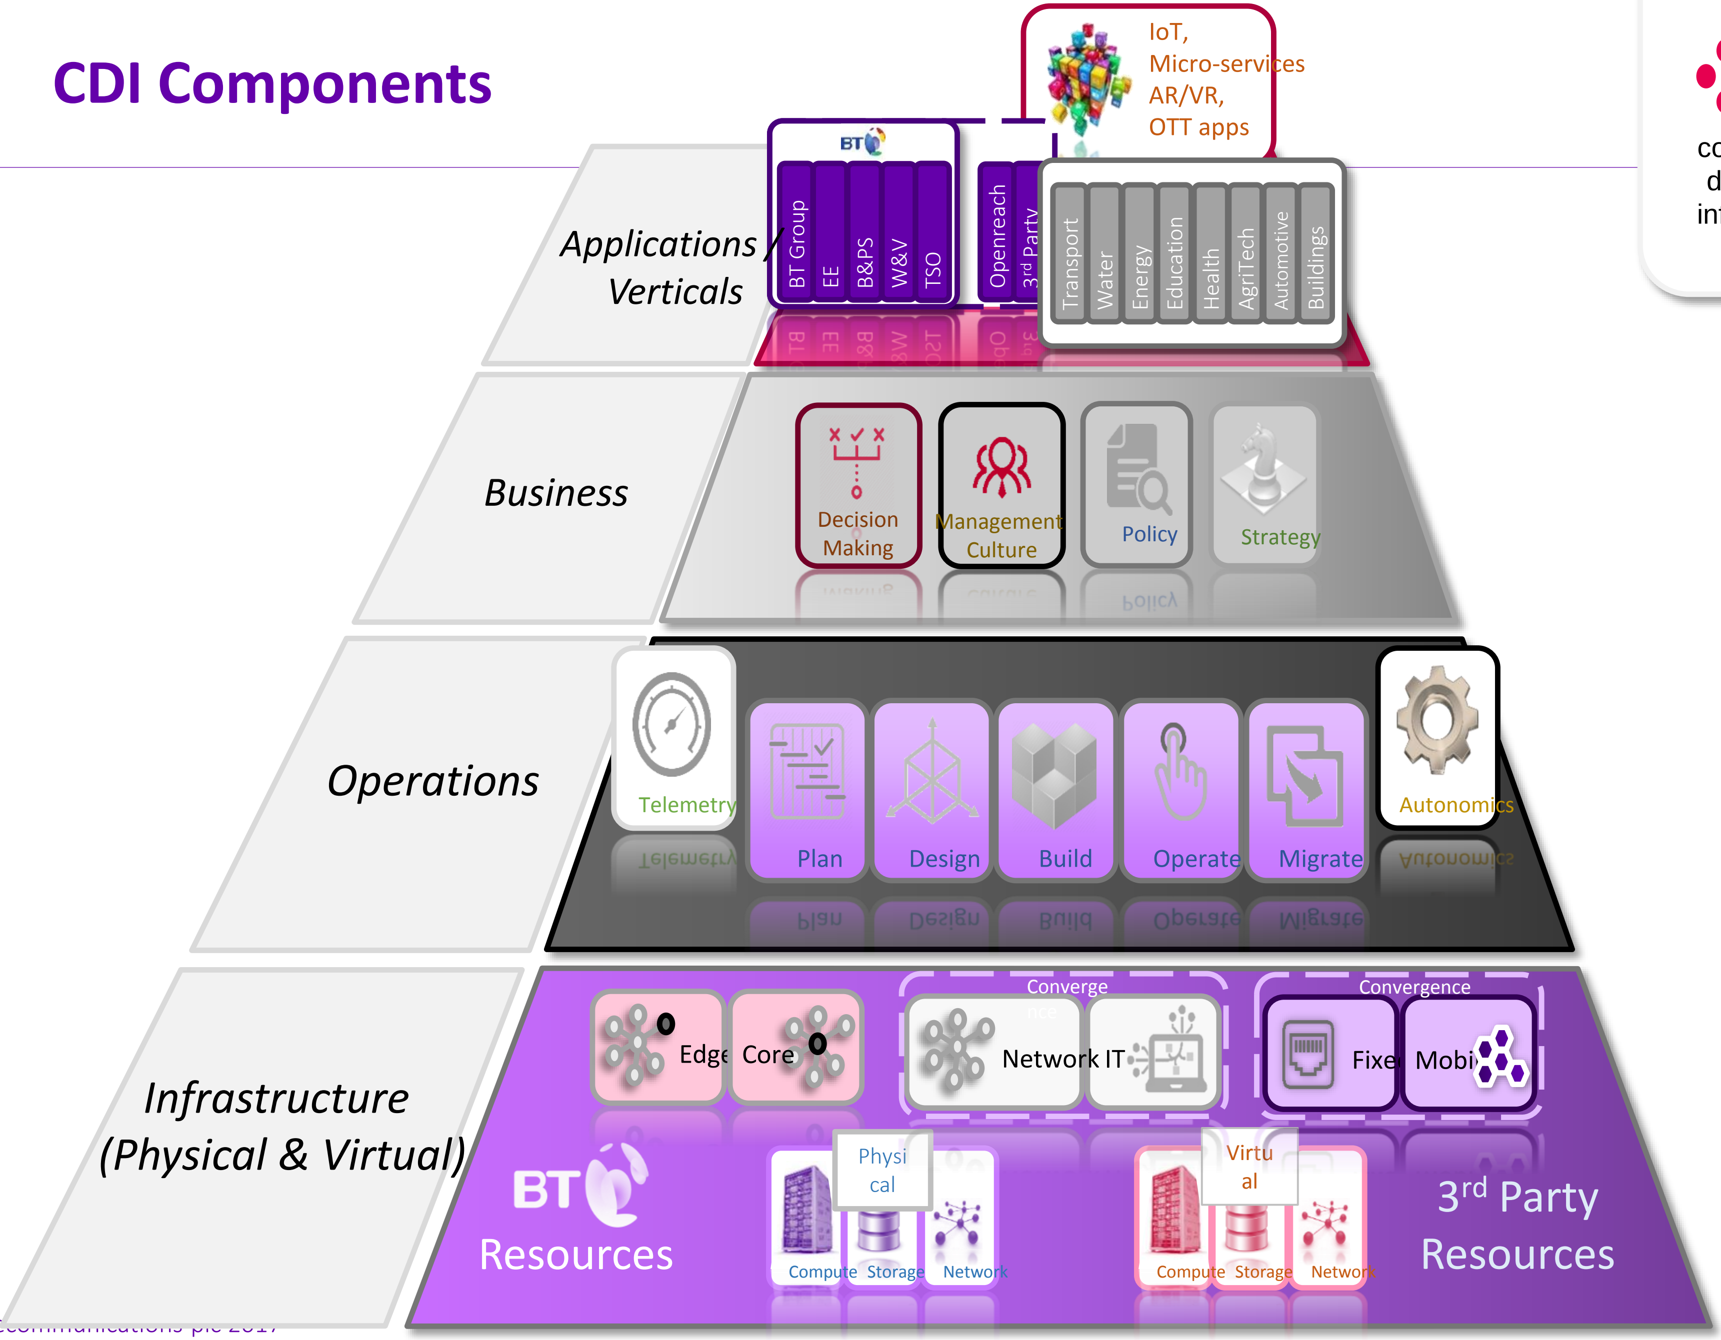

converged digital infrastructure

How do we go about it ?

Collaboration pt 1
Research questions
Where does it sit: the relationship with “AI”
The process of autonomic system design
Autonomic system loops
Special terms
“AI” constraints
Business Autonomous systems
Collaboration pt 2
Summary

Business Autonomics Summary

Domain knowledge and analysis
Machine learning

Full system design that delivers self-CHOP(E)
Double loops
Outer loop – continuous external exploration

Decision control points
Service level vs cost
Stability (risk) vs learning rate

BT