


## Planning, provisioning and operation of virtualised networks

Robin Bailey – Managing Director, ‘Mr STEM’



## Introduction: conventional models over-turned

- Techno-economic modelling for networks is a well-established subject
- A great many tools exist for specific technologies, in addition to the ubiquitous and lowly spreadsheet
- The STEM modelling process helps you capture and portray, in a consistent and transparent manner, all of the critical business drivers of peak or volume demand, capacity, lifetime, locations and engineering effort which link the essential revenue and cost elements of an infrastructure business
- The advent of SDN and NFV has over-turned conventional models, with their focus on capex optimisation, and switched the focus to the opex trade-offs between the known bottlenecks of traditional networks versus the more uncertain but compelling arguments for virtualised networks



## Introduction: ready for business in the cloud

- A very prominent European telco told me recently that their longstanding tools were unable to adapt to the new economics
- In contrast, we have maintained a longstanding principle of technology neutrality which has allowed us to hit the ground running with numerous cloud business models
- The increasing dominance of staff-related opex was anticipated at least five years ago, so end-to-end support for aggregate measures (such as configuration-task hours) is now embedded across our modelling platforms
- We will explore some of the economic realities of migrating to the cloud
- We may reflect that, while once **the network was a service**, now ...
- ... **the data centre is the new network** for the purposes of economic modelling and business insights

# Outline

## The value-chain in a data centre

Migration to centralised facilities

Revenue/cost modelling from hardware and site to IaaS, PaaS and SaaS layers

Build from products or components

## Fulfilment and assurance

Quantifying the scope of tasks and required skillsets (resource mapping)

Sites for hardware tasks vs virtual

Virtual network functions vs hardware

## Cost allocation and pricing

Routine calculation of unit costs

Cost breakdown by component and target-to-cost approach

Price optimisation and sensitivities

## Applications

Making the case for hybrid cloud in the business-plan for a web-scale enterprise

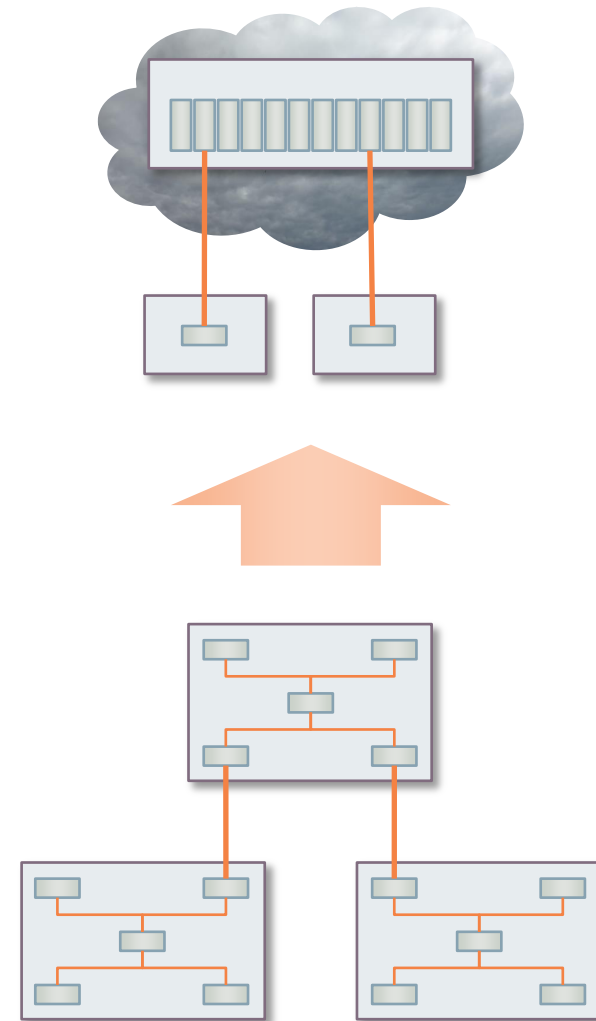
Net benefits for transitioning to SDN

Arbitrage between cloud operators

## The value-chain in a data centre

# Network economics increasingly focused on the data centre

- The data centre is now ubiquitous in carrier and web-scale business infrastructures
- Very high speed fibre and all-purpose IP networking have driven an inexorable migration from local offices to centralised facilities:
  - greater equipment utilisation
  - massive operational efficiency
  - consistent platforms
  - faster response to configuration faults in higher layers
- The data centre is **the new network** for the purposes of economic modelling and business insights



# From hardware and site to IaaS, PaaS and SaaS layers

- Running your own hardware may not be in vogue ...
  - ... but you need a model if you are going to evaluate the alternatives
  - So we have modelled the data centre as a value chain to help determine at which layer you have the best scale to operate
  - The cost of the raw physical assets may be compared with the price of consuming IaaS, and higher layers in turn
- **Physical layer:**
    - compute, storage and network
    - building, power, cooling, UPS
    - deployment, operations, security
  - **IaaS:**
    - $VM_{large}$ ,  $VM_{medium}$ ,  $VM_{small}$
    - $VM_{image}$ , storage, VLAN, firewall, load balance, IP address, etc.
  - **PaaS:**
    - OS, scripting, web, database, etc.,
    - running on elements of IaaS
  - **SaaS:**
    - managed email, backup, etc.
    - running on elements of PaaS




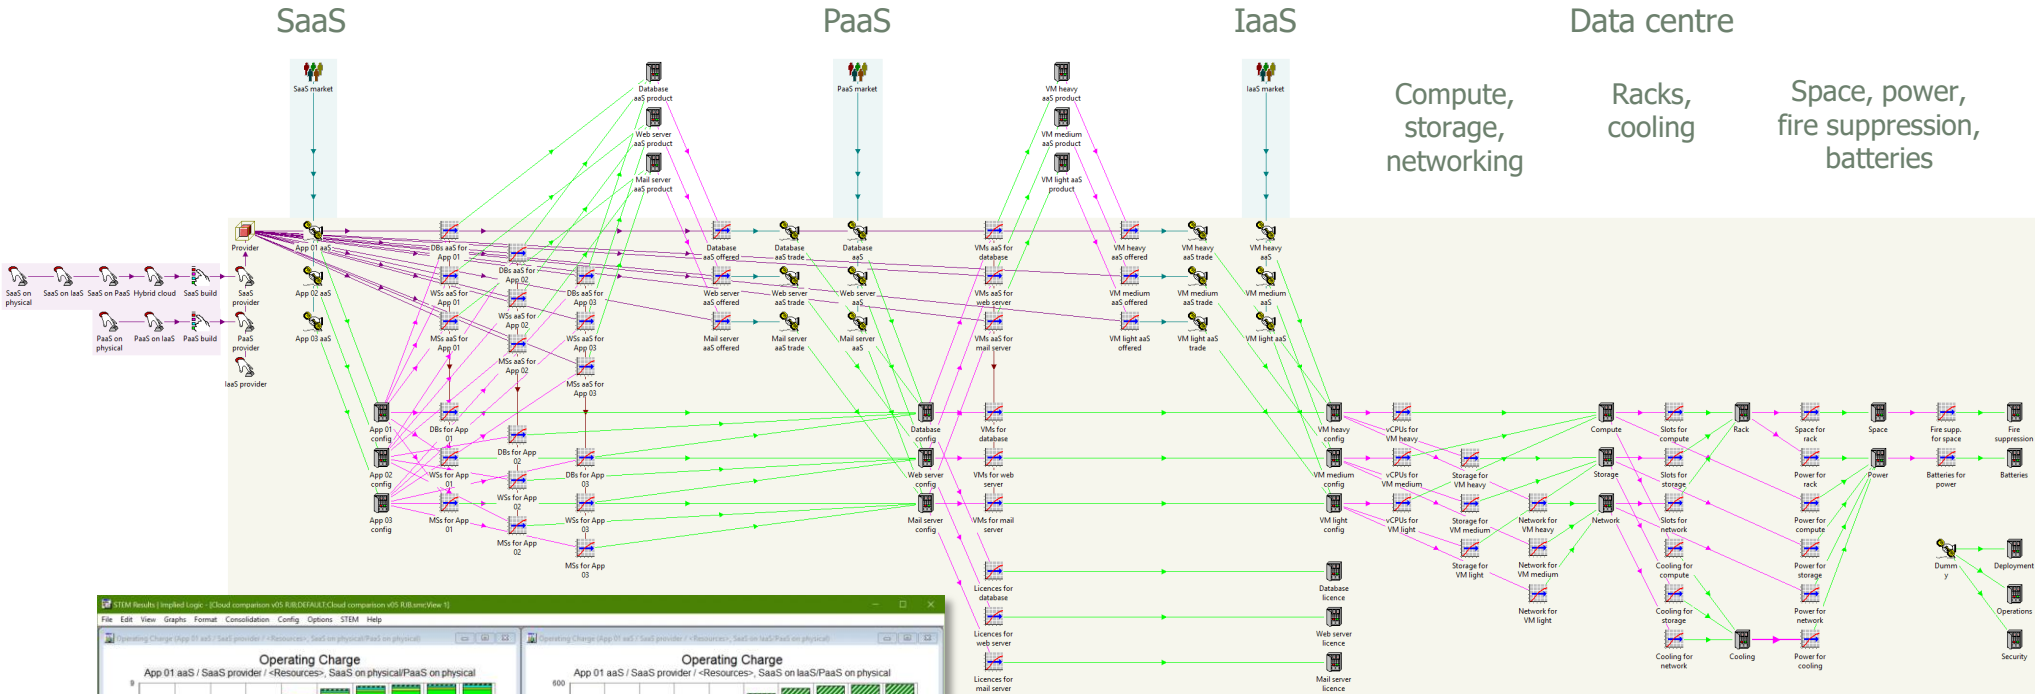
## Build from products or components

- Interesting question as to whether the economics are better for a platform or software operator to build their solution on IaaS (or PaaS) rather than the raw hardware:
  - if an IaaS operator prices keenly, then their margin *should* match the operational benefit for a client not having to maintain its own physical assets and being able to focus on its main value-add
  - however, the initial hype may have allowed operators to over price (they might say, “to recoup R&D”)
  - a fair deal should be negotiated
  - a model is required to establish an objective reference point on pricing
- Our model uses scenarios to compare the following approaches:
  - SaaS on PaaS
  - SaaS on IaaS
  - SaaS on physical
  - PaaS on IaaS
  - PaaS on physical
- You can use the same structure to evaluate which model to use:
  - SaaS vs PaaS vs IaaS vs physical
  - where do you have expertise/scale?
  - what is your value-add/focus?
- **Understanding the value-chain of a data centre is fundamental**





# Physical and aaS technical architectures compared



- SaaS, PaaS and IaaS operators modelled in parallel
- Build from own components or platform/infrastructure *aaS*
- Compare own data centre costs with mix of *aaS* product costs

## Fulfilment and assurance

# Quantifying the scope of tasks and required skillsets


- A key part of the virtualisation story is automation, i.e.:
  - the ability to commission and maintain services without recourse to an underlying hardware layer
  - the potential to turn-up capacity on-demand with no manual input
- For a business case to reflect these operational benefits, it is essential to capture the effort associated with the full range of fulfilment and assurance tasks:
  - first in a conventional network hardware configuration, and then
  - in one or more virtualised scenarios for comparison of cashflow, profitability and NPV
- As well as the overall effort per task, per new/existing customer ...
- ... a mapping across specific roles will yield more detailed insights:
  - customer-service agent, network engineer, data-centre IT support, virtualisation architect, network security consultant
  - team lead, senior management
  - facility services and security
- It is not just the hours, but the hourly rate for each skillset
- It is essential to calibrate a model with this level of detail, and to:
  - understand drivers/dependencies
  - consider minimum staffing levels

## Sites for hardware tasks vs virtual = hidden cost

- If you have network intelligence configured in hardware across a distributed network ...
- ... then you must have engineers at many sites to configure it
- If you consolidate these scattered platforms to a small number of data centres, then:
  - you can expect to achieve higher utilisation in the equipment, but
  - you still need people at each site
  - moreover, for practical reasons, these few sites are likely to be spread out across the territory
  - so separate teams remain a necessity
- In a virtualised architecture, you may need some core expertise to maintain 'utility IP connectivity' at each site ...
- ... but any virtual service may be fulfilled and monitored from a single, central NOC\*
- This critical 'sites assumption' and very real overhead factor for operational resourcing is easily overlooked ...
- ... especially when considering the business plan for a new entrant
- Thus a virtualised future may lower barriers to entry and lead to increased competition

\* See illustration on following slide

A virtual service may be fulfilled and monitored from a single, central *network operations centre (NOC)*\*



## Virtual network functions vs hardware; operational efficiency

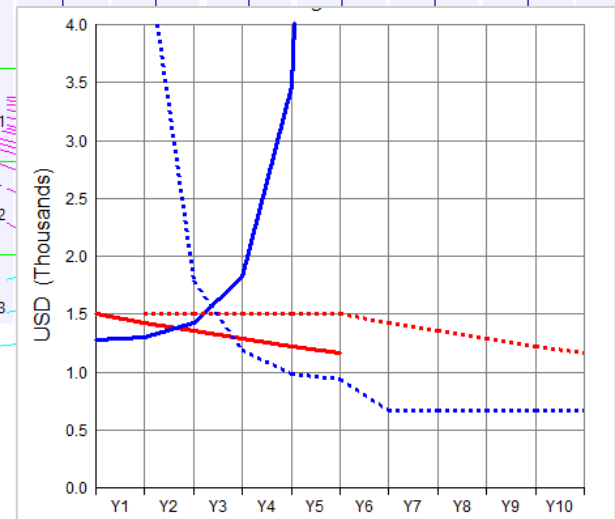
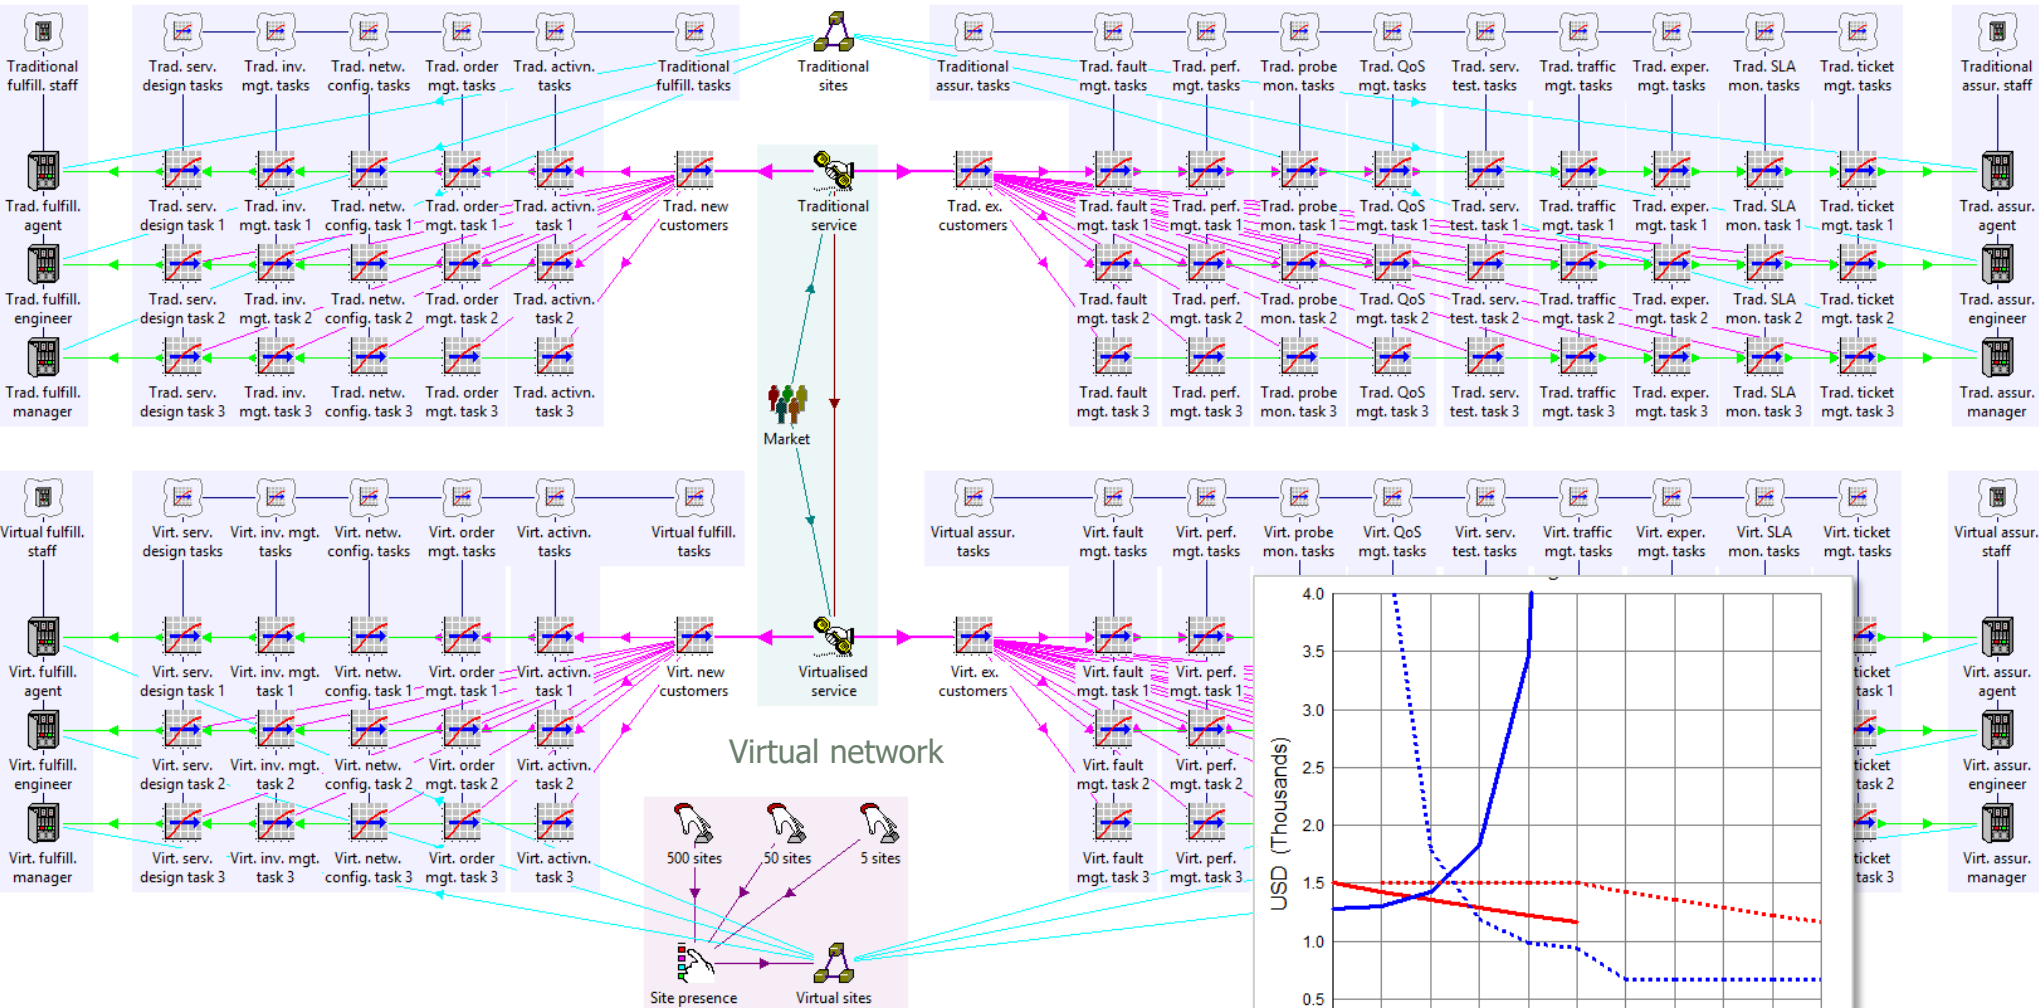
- It is entirely practical to model the economics of virtualised networks, including cost and tech. trends
- Traditional hardware elements still exist, but with standardised, low-complexity configurations.
- High speed networks can bring all the intelligence to the data centre:
  - compute, storage, networking
  - per service/server licences
  - orchestration platform licence
- Less focus on capital efficiency:
  - most of the long-term investment is in the safe bet of the data centre
  - much of the virtual platform presents as an opex item
- As the network becomes more of a commodity ...
- ... and with an increase in white-box solutions and open source ...
- ... the human operational costs will become more and more dominant in future business plans
- Whereas such costs are hard to estimate and were often estimated in the past as a % of capital ...
- ... we anticipate an increasing interest in and demand for well thought-out and consistent approaches to opex modelling

# Labour costs per new customer and per existing customer

Fulfillment tasks

Traditional network

Assurance tasks



Revenue and cost per customer compared, traditional versus virtual

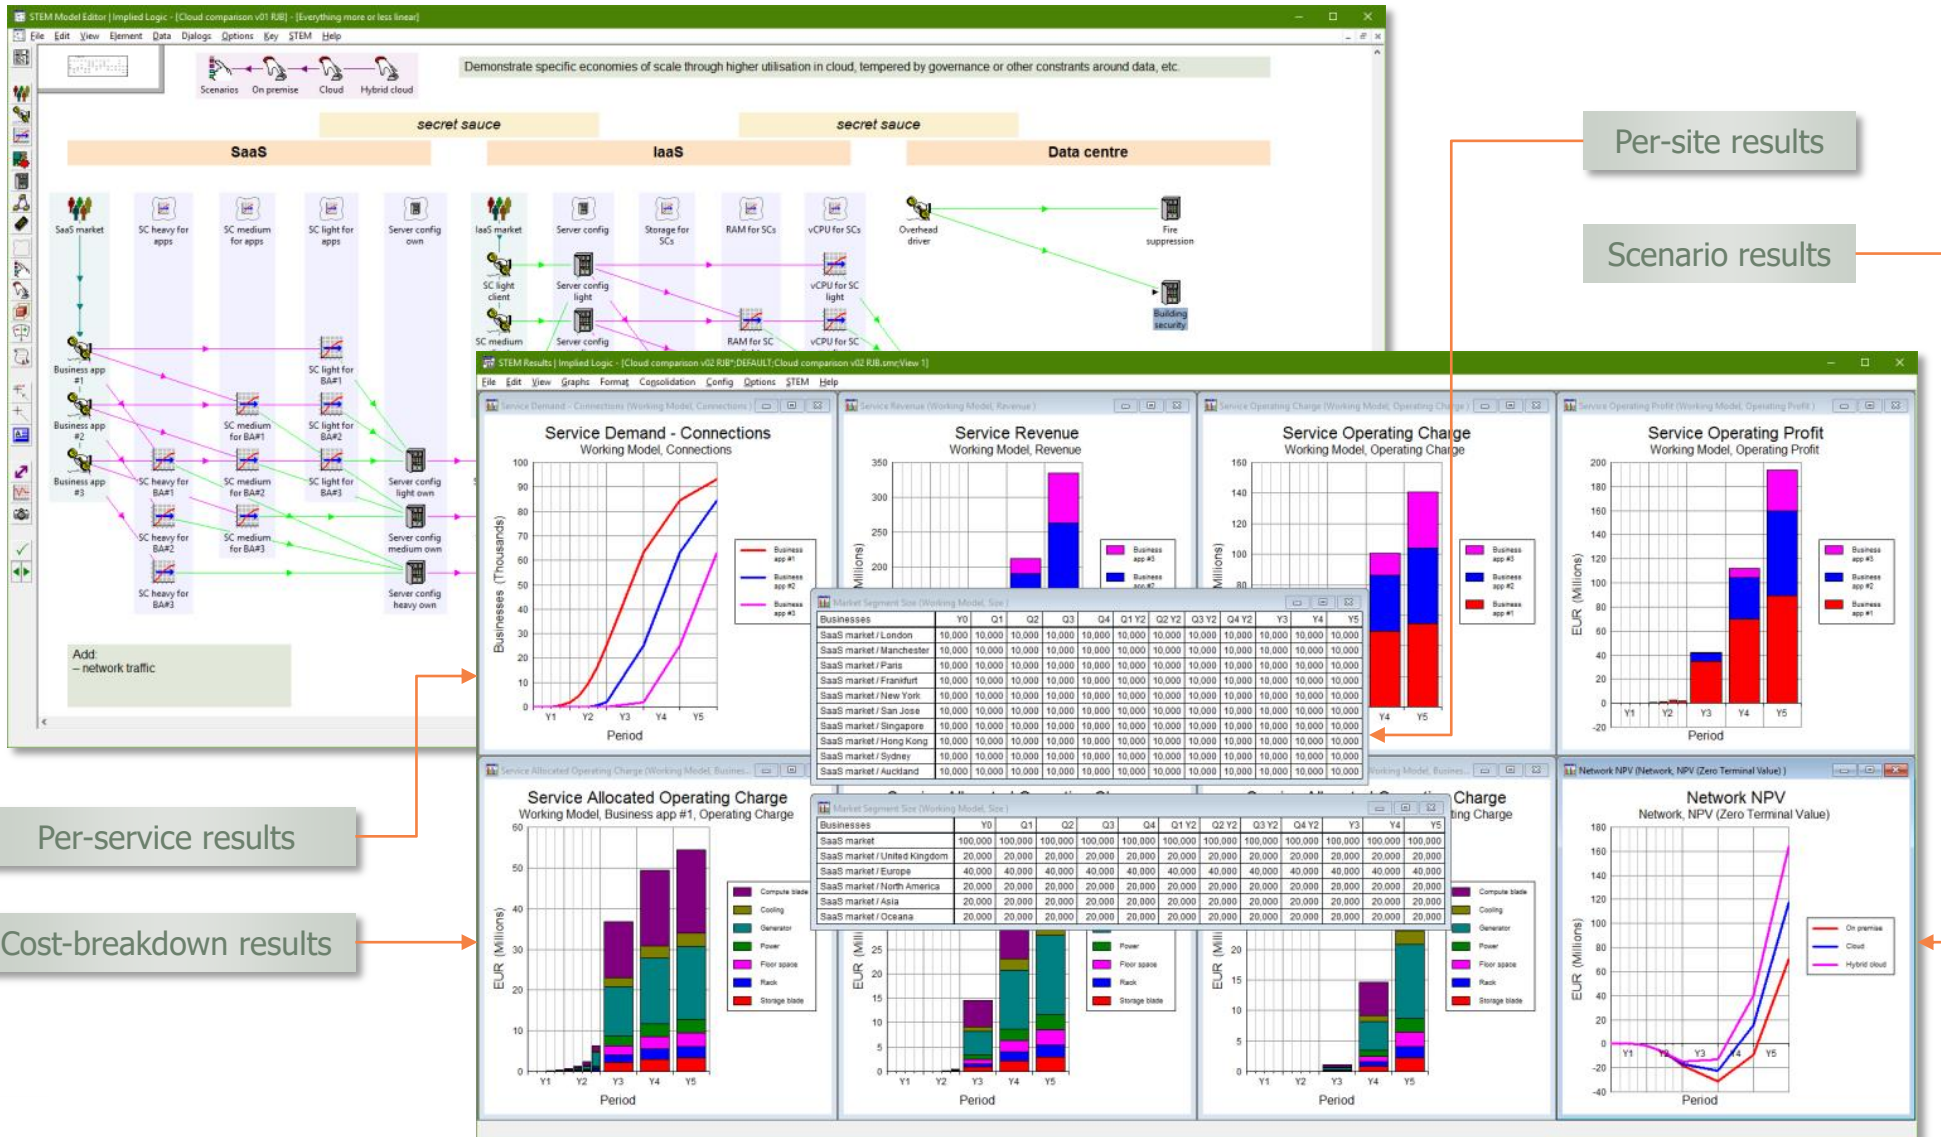
## Cost allocation and pricing



## Routine calculation of unit costs is an essential health check

- It is a fact that many network businesses have developed on the strength of likely revenue vs projected investment cost alone
- Little or no detailed analysis is done of the incremental cost of service provision
- Prices are often set at market rates without reference to whether this is good or bad for business
- Spreadsheet provisioning models have no functionality to allocate costs, nor are their engineering owners necessarily interested in this perspective
- In contrast, our model allocates all direct and overhead resources costs to services 'out of the box'
- The costs of shared resources are allocated in proportion to demand by default, or you can choose:
  - by service volume, or
  - by revenue
- As well as fully-allocated costs, our model also reports:
  - the cost of used equipment, and
  - the direct cost of all resources allowing for the utilisation of intermediate drivers
- This enables you to differentiate between **incrementally profitable** and **overall profitable** pricing

# Intrinsically consistent results across sites and scenarios



## Cost breakdown by component, and target-to-cost approach

- In addition, our model provides a detailed breakdown, by originating resource, of allocated service costs:
  - for every individual cost driver, or
  - to any level of divisional layers (such as SaaS/PaaS/IaaS/physical)
  - opaque or transparent reporting through intermediate service layers
- This capability is vital for assessing the top cost drivers, and also:
  - provides a critical sanity-check for the overall results (compared to the allocated totals alone)
  - can help debug stray (accidental) cost drivers in a model
  - *is terribly difficult to do reliably or consistently in a spreadsheet!*
- *A target cost is the maximum cost that can be incurred on a product (component) and, with it, still earn the required profit margin from a product at a given selling price\**
- In other words, it is ‘a reverse allocation from services of allowed budget for costs’
- Our existing allocation capability enables the following approach:
  - determine credible market rates for service revenue (per customer)
  - subtract required profit margin
  - match remaining budget to cover all allocated costs

\* From *Wikipedia*, the free encyclopaedia: [https://en.wikipedia.org/wiki/Target\\_costing](https://en.wikipedia.org/wiki/Target_costing)

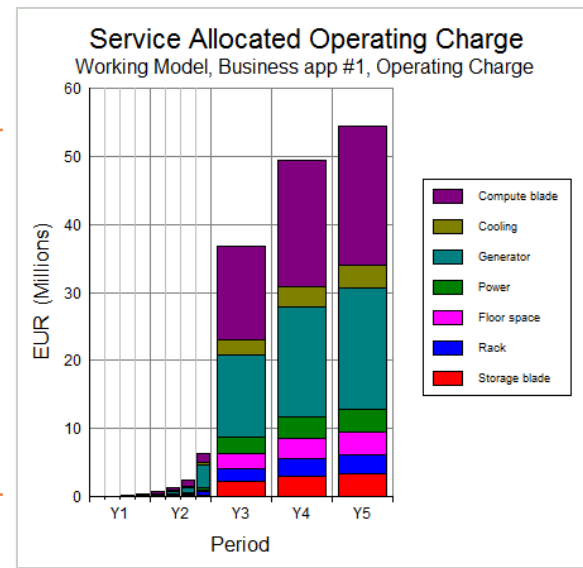
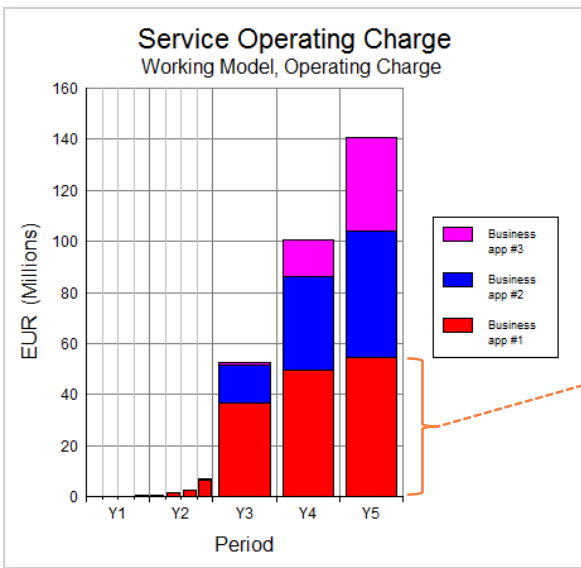
# Every STEM model features cost allocation 'out of the box'

## Cost allocation (to services)

- STEM allocates all resource costs
- Shared between multiple drivers:
  - *pro rata*, by default, but
  - can be overridden; e.g., by revenue
- Aggregates along demand chains
- Provides total allocated cost

## Cost breakdown (by resource)

- Optional breakdown of allocated cost by contributing resources
  - also useful as a debugging tool
- Intermediate services can be used to represent divisional layers (such as *SaaS / PaaS / IaaS / Physical*):
  - *opaque* or *transparent* reporting




## Price optimisation and sensitivities

- A detailed scenario model is the best way to conduct an orderly what-if analysis on an uncertain pricing strategy:
  - explore different segmentations
  - upfront vs recurring pricing
  - NPV impact of incrementally profitable vs overall profitable
  - measure the impact of unforeseen incremental network costs
  - measure the impact of predictable competitive pressure on pricing
- This can be tied in with scenario planning of network design for a truly integrated approach
- Orderly what-if analysis of *the devil you know* should always be balanced with unbiased sensitivity analysis of *the devil you don't*
- Our model includes an integrated point-and-click sensitivity tool which can produce tornado charts for multiple scenarios in seconds
- It can also be driven by third party *Monte Carlo walk* add-ins for Excel
- Or you may prefer to identify unforeseen opportunities and risks in a workshop with trusted colleagues (or, alternatively, with strangers with no preconceptions!)




# Applications



# The case for hybrid cloud in a web-scale enterprise

- Public cloud services offer great efficiencies through massive scale
- They can also offer on-demand flexibility at a better price than owning your own compute assets which would be mostly dormant\*
- Private cloud capabilities are less efficient and may be regarded as a non-core business distraction, but may have to be part of the mix:
  - due to regulations on sensitive customer or financial data
  - due to cross-border privacy issues
  - if custom processing is required
- How will you determine the best approach?
- The data-centre value chain is just right for comparing the options:
  - the public cloud approach may be characterised by the consumption of the relevant SaaS, PaaS and IaaS services (mix of fixed and variable) with the consumed service revenues as the effective cost
  - the private cloud approach takes the same requirements and pushes them through to the underlying hardware costs
  - or you may consider a suitable public-private cloud split
- **Whatever the technical options, a rational investor will increasingly demand to see the underlying economics!**

# On-demand processing can be shared in a public cloud





# Demonstrating the net benefits of transitioning to SDN



- This conference knows all about business models for traditional networks
- By modelling the data centre, in conjunction with the associated operational impacts, it is entirely feasible to evaluate the virtual network alternative:
  - cost of deploying new platform
  - uncertain learning curve for new mode of operation
  - migration to white-box hardware
  - conjectured reduced operational effort and associated dollar cost
- If you are really careful and thorough you can try to calculate the absolute cost of each option
- However, if your objective is to determine which option is better, then it suffices to model what is different and look at an **NPV delta**
- In practice there may be multiple dimensions to the decision making:
  - what to do if such a market or technology eventuality occurs?
  - which option is better in market A compared to market B?
  - is first mover a curse or advantage?
- **Never has network economics been more relevant!**

# Identifying the scope for arbitrage between cloud operators

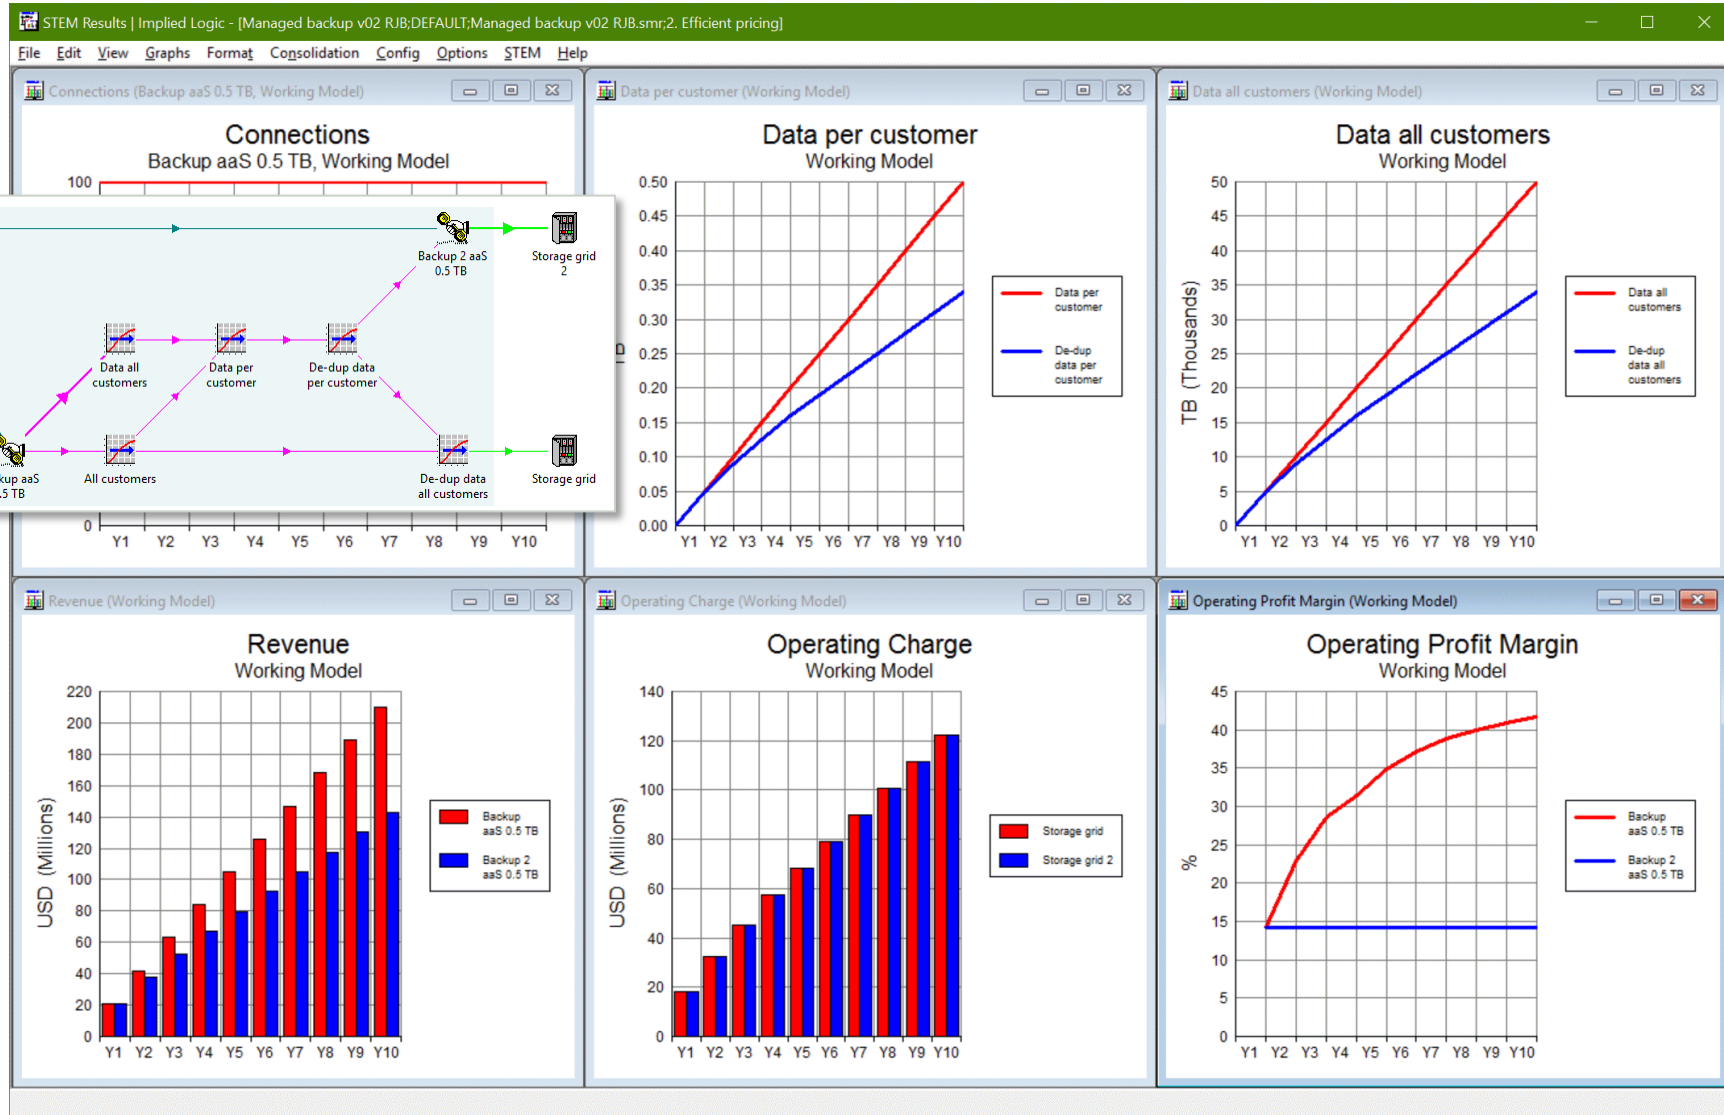
- Consider a backup platform which de-duplicates data as it enters the grid to minimise storage usage
- This capability is commonly sold as a cloud backup service, either:
  - A. priced per GB assured, or
  - B. priced by actual consumption
- The data efficiency of the solution typically increases with usage, but to whose benefit?
  - A. service provider margin actually increases with usage (*you thought your cloud had a silver lining?!*)
  - B. service provider revenue increases on a flat margin\*
- It is very easy to model these two pricing strategies as scenarios:
  - *Option A* may well show a higher revenue if there is no competition
  - *Option B* is likely to be profitable too, or can be structured to cover costs 'sooner rather than later'
- This may demonstrate a classic *arbitrage* where a new entrant could challenge existing operators with a more efficient offering (*and have the incumbent's breakfast!*)
- **Much of the current excitement about cloud is technical, but the fate of businesses will still be determined by economic realities**

\* See illustration on following slide

# Service-provider margin increases with usage, or revenue increases on a flat margin



# Pricing per GB assured = opportunity for efficient pricing



# Conclusions

- Network business models were always rich in dimensions:
  - customer, service, location
  - timing, technology, vendor
- Tremendous complexity even when the dominant hardware cost drivers were clearly parameterised
- Onus on modelling increases as results become more volatile due to the less predictable dynamics of human resourcing
- Virtualisation creates an economic landscape where most of the traditional complexity of sites and geography is removed ...
- ... but these aspects remain relevant while the business case is quantified and proven
- There was already a compelling case for the use of tools which handle these complexities
- Tools have evolved to support the most topical, flow-related aspects of service provisioning
- We have considered just a few specific scenarios which can be illuminated by diligent modelling
- We anticipate an upswing in such activities for the many commercial possibilities which will emerge as SDN/NFV enter the mainstream

# Annual STEM User Group Meeting

- Wednesday–Thursday 05–06 October 2016, King’s College, Cambridge, UK
- Interactive sessions on business planning for convergent services and product-profitability analysis
- Master classes for established users in parallel with fast-track training for newcomers
- Guest presentations from operator and vendor clients



## Robin Bailey

Managing Director, 'Mr STEM'

Mobile: +44 7776 198458

[robin.bailey@impliedlogic.com](mailto:robin.bailey@impliedlogic.com)

*Helping great minds to think more clearly about business logic*

