PSIRP
Publish-Subscribe Internet Routing Paradigm
FP7-INFSO-IST-216173

DELIVERABLE D4.6

Final Evaluation Report on Deployment Incentives and Business Models

Title of Contract	Publish-Subscribe Internet Routing Paradigm
Acronym	PSIRP
Contract Number	FP7-INFSO-IST 216173
Start date of the project	1.1.2008
Duration	33 months, until 30.9.2010
Document Title:	Final Evaluation Report on Deployment Incentives and Business Models
Date of preparation	4.5.2010
Author(s)	Jane Tateson, Trevor Burbridge, Dirk Trossen, Mark Ain (editor)
Responsible of the deliverable	Jane Tateson
	Phone: +44 1473 644820
Email	jane.tateson@bt.com
Reviewed by	Dirk Trossen, Mark Ain, Trevor Burbridge
Target Dissemination Level	Confidential
Status of the Document	Completed
Version	1.0
Document location	
Project web site	http://www.psirp.org/

SEVENTH FRAMEWORK PROGRAMME

European Commission
Information Society and Media

Table of Contents

This document has been produced in the context of the PSIRP Project. The PSIRP Project is part of the European Community's Seventh Framework Program for research and is as such funded by the European Commission.
All information in this document is provided "as is" and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.
For the avoidance of all doubts, the European Commission has no liability in respect of this document, which is merely representing the authors view.

Executive summary

This document takes the PSIRP architecture as its starting point, and distils from it a set of key properties which are enablers for new applications and services. These enablers create a new type of information networking, which we call the Information Cloud.

The key properties are:

- The Information Cloud decouples information from infrastructure and location.
- Information is organised, accessed and routed according to its meaning.
- The content of information is separated from its governance, which enables dynamic loosely-coupled relationships to flourish.
- Information can be tagged in multiple ways, providing rich meta-data.

Working from the premise of these capabilities, three sectors were identified as being the strongest contenders for benefiting from what the Information Cloud has to offer. They are Government, Collaborative Business ICT Services and Content-Centric. Primary benefits to the government are identified as being around cost-effectiveness, accountability and citizen-centricity. Primary benefits to collaborative business concern simplifying multi-service business, enabling dynamic supply chains and de-risking expansion into new markets. Primary benefits to the content-centric sector revolve around media rights management, providing assurance of authenticity, democratising media generation, and enabling multimedia immersive environments.

Having developed stories for the business opportunities in these sectors, the stories are mapped to their dependence on technological components, in order to begin to understand how migration to the Information Cloud would take place. The drivers for network transformation are then built into a system dynamics model of the economy of the network, in which application sectors generate demand for Information Cloud networking. In this work we distinguish between two Information Cloud solutions: the first is a ‘shim layer’ approach which represents an information-centric overlay on to an underlying IP network; the second solution we refer to as ‘native’ Information Cloud, in which the full PSIRP architecture has been implemented. In so doing, we are able to distinguish between market triggers for an overlay network and market triggers requiring re-engineering of network nodes. We model the ‘native’ solution as being more expensive but also more profitable than the ‘shim layer’ solution. We then run the system dynamics model varying the levels of investment provided by the three application sectors, and also varying the investment thresholds of the network operator. The aim is to identify, for a given level of application sector investment, the appropriate investment strategy for the network operator.

We conclude by observing that the network operator would be well served to invest in developing Information Cloud applications with customers from the three sectors, in order to stimulate demand for the network transformation, and suggest a sensible order of priority for this. We also observe that Packet Level Authentication forms a key component of the capabilities of the Information Cloud and deserves further technological and business analysis, and suggest that it would be a key enabler for media rights management, which is likely to be a powerful business driver for the Information Cloud.

1. Introduction

As described in D5.5 [PSI2010], we appreciate that fundamentally altering the Internet paradigm is perceived as a challenge of monumental proportions and, therefore, one that most people would rather not consider. Attempts to modify the underlying Internet architecture have been, up until now, impossible to manifest, due to a lack of:

- **Understanding** (i.e. poor appreciation of what the solution entails)

- **Perceived urgency** (i.e. no recognized impending operational catastrophe)

- **Motivation** (i.e. no perceived benefit, no gap to fill etc.)

Unfortunately, conflicts of interest also arise in commercial and industrial areas because drivers such as profit and politics dominate the allocation of resources which largely shape the future of the Internet.

In this document, an attempt has been made to address these issues. We contribute to improving understanding, by reprising what the PSIRP architecture involves and its key principles. However our main effort is to address in detail levels and types of motivation for achieving information-centric network transformation. We identify the deployment incentives and business models that PSIRP offers, initially from the perspective of application sectors, but ultimately from the perspective of the network operator.

Having outlined the key principles, architectural features and dependencies of PSIRP, we describe three of the most important sectors that we believe will drive the demand for an information-centric internet, or Information Cloud. We identify the drivers for those sectors, highlighting business benefits and opportunities. We then carry out a detailed modelling exercise to explore the relationship between application drivers and their technological dependencies; and to couple these with application sector investment levels and network transformation investment levels, within a simplified economic ecosystem.

We are primarily interested in determining what would be a good strategy for a network operator, in order for them to maximise their profit, in meeting the demand of new products and services in the information-centric sphere. Depending on the assumptions made in the model, in terms of a willingness for application sectors and network owners to invest, we identify the resulting level of network transformation and network traffic volumes and network profits. A conclusion may then be drawn on the urgency for action in transforming the network into an information-centric medium.

1.1 Abbreviations

PSIRP	Publish Subscribe Internet Routing Paradigm
PLA	Packet Level Authentication

2. Principles of PSIRP

2.1 Everything Is Information

The core principle of PSIRP is to create an information-centric Internet, in which everything can be classed as information. The choice of how small or large a piece of information is will be determined by a given application. In principle, a unit of information could be as small as one bit, carrying a binary message, or as large as a document of several megabytes. Even a video stream could be treated as one piece of information, although, owing to the special requirements of video, it is somewhat of a special case, as will be discussed below.

As introduced in D2.2 [PSI2008a], information items can also serve as metadata for other information items, enabling powerful new functionalities. Even executable code can be viewed as information. So the vision of PSIRP goes far beyond merely being able to access facts efficiently.

2.2 Networks of Information

Information-centricity revolves around a new approach that matches publishers of information to their subscribers using a process known as rendezvous. During this process, a piece of information is given a statistically unique label that creates a transient relationship between the publisher and the subscribers to this information, which is called the **rendezvous identifier** (RId).

Information is not intended to exist in a completely flat information space. In order to be able to provide restricted access and governance of information, as well as efficiency of operation, information is grouped into *information networks* using the concept of scopes, as introduced in [PSI2008d]. This is illustrated in Figure 2.1.

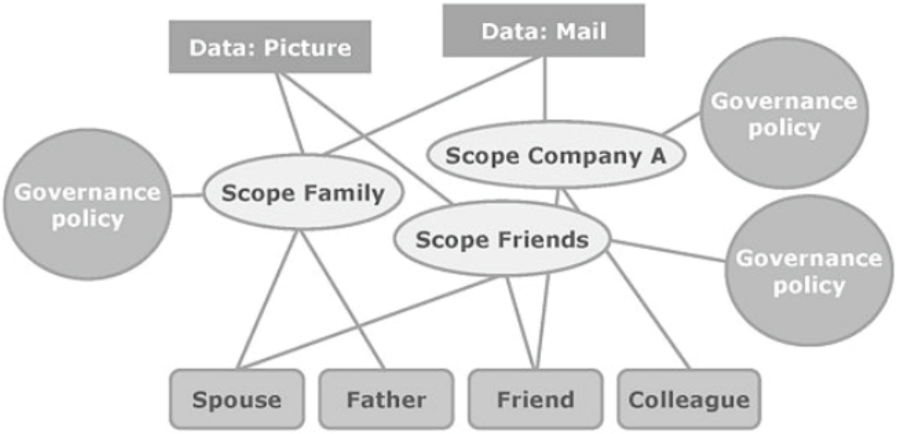

Figure 2.1: Illustration of scopes

Each information network represents information in itself and therefore requires an identifier in order to be addressed. As they represent a special class of information, they have a specific type of identifier, a **scope identifier** (SId) which is a subclass of rendezvous identifiers. Each scope identifier represents what is known as a *rendezvous point* which implements the information network.

2.3 Grouping Information within Networks

Intermediate in scale between individual pieces of information and information networks (scopes) PSIRP also enables *information collections*. These are particularly useful for handling large pieces of content such as video streams, for example. In a video stream, the overall bucket of data is large and belongs semantically together. However users may wish to access only part of the video, or even an individual image. This requires the smaller elements of that stream (frames or even images) to be individually labelled, but they must be grouped together for efficient caching and forwarding. This grouping can be carried out by exploiting the use of information items as metadata. The metadata information item must represent the RIds of the individual elements of the video stream that need to be grouped together. Rather than this metadata being a list of RIds, the RIds may be created as **algorithmic rendezvous identifiers**. That is, the RIds of the individual elements may be created algorithmically from the identifier of the metadata item.

2.4 Grouping Networks of Information

We have just seen how *information collections* can be used to group information items together within an information collection. However, it is also possible to collect information networks themselves. These collections of information networks can represent networks within organisations or networks within social structures, e.g. social networking sites. They are grouped together by using multiple successive scope identifiers or through the use of **algorithmic scope identifiers** which encode the collection of information networks (or rendezvous points) (or SIds) which need to be grouped together within a larger scope, such as Facebook, BT, UK Corporate or EU etc. Whilst the algorithmic scope ID may be divulged to untrusted parties, it will not enable reverse engineering of the information network collection. This enables *private information networks* to be generated.

3. Overview of Architecture

3.1 Entities

- *Information Item*: any data can be labelled with an identifier, forming an information item that our architecture can process. Application identifiers are resolved into rendezvous identifiers, either directly by an application or via a helper function.
- *Information network* (or scope): information items can be grouped into networks of information, which are labelled with a scope identifier. An information item can be associated with one or more scopes and these scopes can be interpreted at various levels of abstraction.
- *Information subscribers* and *publishers* that consume and create information items, respectively.
- *Domains*, which are administrative physical network areas that can be connected using the inter-domain forwarding architecture. Propagation of information items within and across domains takes place via labelling routes with forwarding identifiers.

Figure 3.1 gives an example overview of these key entities, based on the current stage of architecture development.

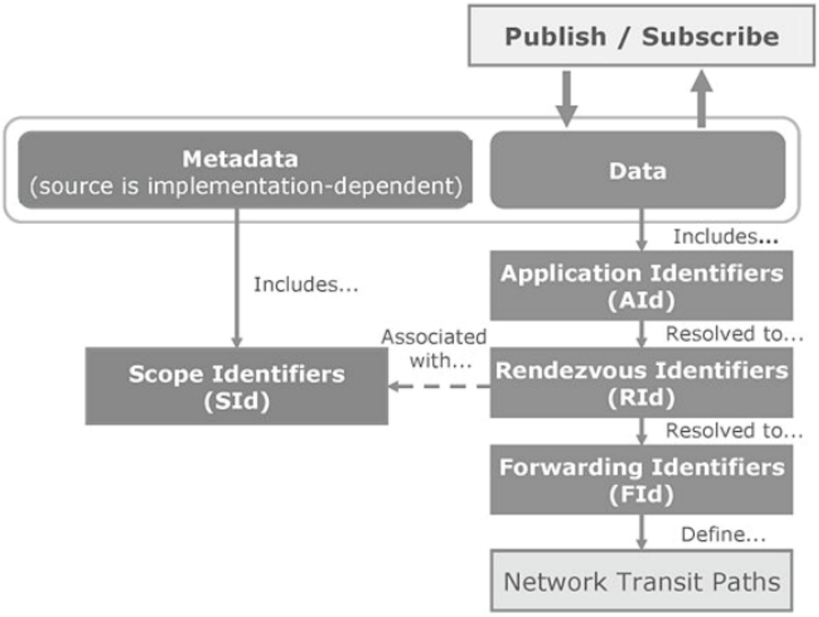

Figure 3.1: Relationship of key entities

3.2 Key Processes

- *Rendezvous* is the process whereby subscribers and publishers can meet within the network. Publishers use rendezvous to determine subscriber information that is later used by the *topology management & formation* function. The scope determines the part of the rendezvous system that is used by the network. Rendezvous can take place at the level of the local node, using the blackboard to manage local publish and subscribe operations (see below). Rendezvous also takes place over local links. In general, however, physical network devices known as *rendezvous nodes* (RNs) host *rendezvous points* (RPs) which are logical, fixed or non-fixed indirection points for publish/subscribe network communication. Rendezvous nodes within a collaborating set of network domains form *rendezvous networks*, where global reachability requires rendezvous network interconnection. Note that the rendezvous system is a policy enforcement point in the architecture and a mechanism for enabling freedom of choice for network end-points, whereby a *rendezvous identifier* is associated with policy-compliant data dissemination graphs for publication delivery.

- *Topology Management & Formation* is responsible for managing intra-domain delivery topologies as well as forming inter-domain forwarding graphs for publications. This process derives sets of *forwarding identifiers* that are used by the *forwarding* function. Note that PSIRP assumes a structure of autonomous systems interconnected via inter-domain mechanisms, similar to the current Internet, where each system is assumed to be individually governed by an organization to optimize its use of resources.
- *Forwarding* pertains to data delivery within a single administrative domain or across multiple domains. Its function is to use the information provided by the rendezvous and topology functions to make decisions on forwarding nodes.
- *Helper functions* are used to extend the core functionality of the network architecture and provide services in relation to functions such as management and transport. Examples would be network management solutions (such as SNMP), diagnostic tools (ping, traceroute, tcpdump), caches and proxies etc.
- *Network attachment* is responsible for discovering network attachment points and configuring components in such a way that communication becomes possible. This includes setting up, maintaining and renewing state associated with local link configuration and authentication of nodes and users.

3.3 Service Model and API

There are currently four classes of network service in the architecture, which are described in detail in D2.3 [PSI2008b]. Here they are simply listed for completeness:

1. A low-level model that exposes the network forwarding and rendezvous/topology formation functions.
   a. Forwarding
   b. Rendezvous
   c. Publish & Subscribe
2. A mid-level memory object model for sharing memory objects, through the network, from publishers of new (or revised) objects to subscribers.
   a. Simple unreliable page service
   b. Simple memory-object service
3. A mid-level channel model to provide channel services such as flow or error control.
4. Higher-level service models
   a. Many-to-one services
   b. Concast-like services

3.4 Component Wheel

The PSIRP conceptual architecture is based on a modular and extensible core, called the **PSIRP component wheel**. The architecture does not have a traditional layered stack, but rather components can be decoupled in space, time and context. The novelty of PSIRP is to use a publish/subscribe style interaction throughout the conceptual architecture and thus support a modular protocol organization.

The node architecture is structured around the component wheel. Applications, components, and the network and disk interfaces are all attached to the local blackboard (BB), at the centre of the component wheel. The blackboard is used to share publications, both data and metadata, and the publish/subscribe paradigm is used to signal changes to publications.

All system metadata are potentially visible to all components through the blackboard. Hence it is easy to add a new component into the system without having to change the basic system calls and the API.

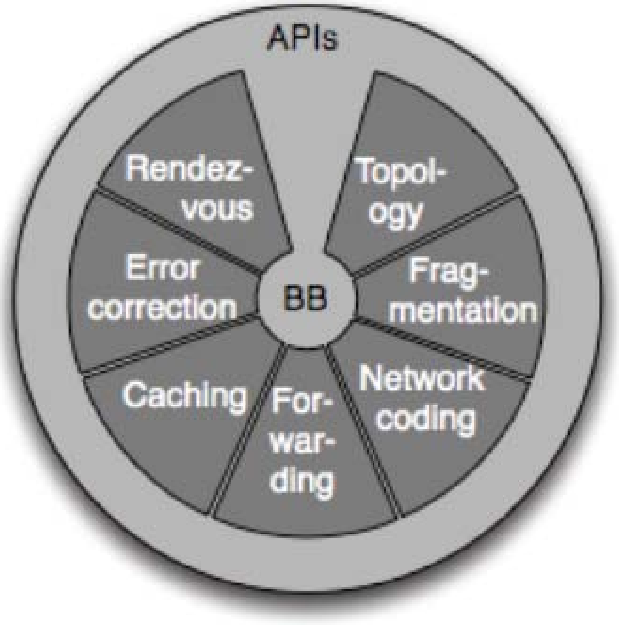

Figure 3.2: PSIRP component wheel

When a publish operation is invoked by an application, memory for the publication is allocated and subsequently the publication is placed on the local blackboard.

A local subscriber can subscribe to receive a notification of when a certain event occurs. For example, if a component wants to know when a new publication is added in the local blackboard, it subscribes to notifications about updates to all locally known scopes. As all publications belong to one or more scopes, including the scopes themselves, addition of new scopes can also be detected.

We see, once more, that the publish/subscribe mechanism operates at all levels within the architecture.

4. Security Considerations of PSIRP

Whilst for some architectures and services, aspects associated with security are essentially 'add-ons', that is not true in realising the capabilities of the PSIRP architecture. Much of the functionality that would be most attractive to end-users of a PSIRP architecture, relies upon certain security considerations, particularly authentication. Therefore there is a brief overview of security considerations given here, and an introduction to Packet Level Authentication, which we believe is critical to many of the assumptions made about what PSIRP could deliver.

The security goals of PSIRP include confidentiality, integrity, availability and accountability. Solutions to these goals are discussed in the architecture document [D2.3] primarily from a network level perspective.

4.1 Direct and Indirect Cryptographic Associations

Direct cryptographic associations are used with fixed data objects, while indirect cryptographic associations are used with data channels, such as a voice call. In the direct association case, there is a direct mapping between RIds and content, by for example, calculating one-way hash values over a block of data. For real-time media we can use indirect cryptographic associations, like public-key techniques. A public key can be used to form part of an RId or SId. The trust is based on the association between a public-key pair and a legal entity, such as a person or company. Unlike the case of direct cryptographic association, the identifier does not need to be changed if the content is modified.

4.2 Information Freshness

With the same content stored at multiple caches, version control becomes necessary. The expected approach is that an SId and the related public-key pair can be used to associate a set of different versions of a document, using algorithmic identifiers. An algorithmic identifier denotes that the subsequent identifiers can be associated together in a provable way, e.g. Lamport's one-way hash chains.

4.3 Rendezvous Mechanism Based on Direct Cryptographic Association

For direct cryptographic association, a trust relationship between content subscriber and publisher is not necessary – the subscriber can verify the integrity of the original content. However, this approach means that the RId becomes critical and so must be delivered over an integrity-protected channel between publisher and rendezvous system and between rendezvous system and subscriber.

4.4 Rendezvous Mechanism Based on Indirect Cryptographic Association

When a public-key pair is associated with content, the resolution mechanism is responsible for mapping the subscriptions to the right public keys. This adds an additional indirection layer to the system. Here, the information object has an 'owner' and the subscriber needs to trust the host with the private key which is providing the content.

4.5 Network Packet Authentication and Access Control

Providing authentication enables any legitimate party in the network to test the authenticity of a packet and to drop the packet if it fails those checks. To check authenticity, the network does not have to know the identity of the packet originator. This can protect confidentiality while allowing certain forms of accountability, such as the packet coming from a user within an organisation.

4.6 Packet Level Authentication (PLA)

Packet Level Authentication (PLA) [Lag2008, Can2005] assumes that per packet public-key cryptographic operations are possible at wire speed because of new cryptographic algorithms (elliptic curve cryptography [Kob1987, Mil1985]) and advances in semiconductor technology [Jar2007, For2008]. This means that every node in the network is capable of checking the authenticity and integrity of packets and hence is capable of immediately detecting and preventing attacks.

PLA adds a separate PLA header to the packet containing a Trusted Third Party (TTP) certificate. This guarantees that the sender is a valid entity, authorized by a trusted third party. Furthermore, PLA uses identity-based implicitly certified cryptographic keys [Bru2007] which means that the sender's public key can be calculated from the TTP certificate. This means that there will be sender accountability, which has many benefits, as we will see later. It is also envisaged that the PLA's TTP certificate will be used to contain the public key of a PSIRP scope. This certificate will be given to authenticate subscribers and so ensures that only they are given access to a particular scope.

5. Information Cloud

In working to identify the new market opportunities that the PSIRP architecture enables, the term 'Information Cloud' has been coined. The Information Cloud builds upon the PSIRP architecture to understand the capabilities that this architecture offers from the perspective of a range of different types of business. First we consider four key features that PSIRP enables, which become the four key features of the Information Cloud:

1. The Information Cloud decouples information from infrastructure and location.
2. Information is organised, accessed and routed according to its meaning.
3. The content of information is separated from its governance, which enables dynamic loosely-coupled relationships to flourish.
4. Information can be tagged in multiple ways, providing rich meta-data.

5.1 Decoupling Information from Infrastructure and Location

The Information Cloud creates a highly distributed cloud of information items: large and small, which are decoupled from infrastructure and information location. This means that applications can be written purely in information-centric terms.

5.2 Organising, Accessing and Routing Information According to its Meaning

Organising, accessing and routing information according to its meaning catalyzes the proliferation and caching of information items. In the same way that peer-to-peer networks operate, in which only the requested content is important, not its location, the distribution of content through the network will naturally respond to demand. Copies of content that is heavily requested will be able to proliferate and be cached close to the users requesting that content.

These two features mean that, as the content is close, in terms of latency, to where it is needed, application performance will be optimised and the use of the infrastructure will be optimised.

5.3 Separation of Content from Governance

The PSIRP architecture enables governance of information items to be independent of the content of that information. Governance will typically be directly associated with scope IDs, not with the information items that lie within that scope. This is important as it creates a step change in the 'fluidity' with which information relationships can be created and altered. Firstly, information items can be moved in and out of scope rapidly. Secondly, the governance of information items, mediated via the scope ID, does not need to know the full content of the information item. This means that the action of governance can be limited to only those aspects of the information that are relevant to it. Remembering that an information item could be something sophisticated, such as the interface to a company's ordering system, this dynamic governance is key to enabling short-term relationships amongst players to be set-up and taken down rapidly. As we will come back to in the section on Collaborative Business ICT Services, separation of content from governance will catalyse a new way for entities to do business.

5.4 Tagging Information for Rich Metadata

Tagging of information is inherent in the PSIRP architecture. Elements in the Information Cloud are associated into scopes, which can have policies associated with them, each expressed as another element in the Information Cloud. This could correspond to who the recipient should be, a time to live for the information or further information that allows authentication and validation of the original information. Tagging of information, therefore, offers application writers the opportunity to add much richer functionality to the process of interacting with that information, which leads to many interesting capabilities, as we shall see later.

However, the action of policies need not be encoded explicitly in associated information elements. Policies can also be expressed in the way that the information is handled as it passes through the network, according to policies defined between publishers and Internet Service Providers. A powerful application of this would be to identify a particular scope with a certain level of priority or quality of service. A node in the network would simply need to recognise the scope ID to filter the packets accordingly. The ease of this transaction, without the need for deep packet inspection, brings the possibility of providing dynamic QoS to end-users, for the first time, and we believe this will be critical to the acceptance of the PSIRP architecture and the revenue of ISPs.

6. Key Opportunity Areas

Three key opportunity areas have been identified, at this stage: Government, Collaborative Business ICT Services and Content-Centric.

6.1 Government

There are many opportunities that the PSIRP architecture, or Information Cloud, can offer to a Government. The following discussion will be mainly from the perspective of the UK, but will apply to most governments.

The Information Cloud offers a range of benefits to Government, in a win-win fashion: helping to lower infrastructure and network costs, as well as increasing functionality and performance. In addition, if the Information Cloud approach is followed to its logical conclusion, there is the opportunity for a political shift towards citizen-centricity.

We will begin by describing the drivers for a PSIRP architecture in the context of Government, for example, cost-effectiveness. We will then describe the functionalities expected from an Information Cloud approach, and finally we will review some more profound benefits.

A typical way for a government to store its citizen data is in complex multi billion pound (Euro) silos, divided up according to government department. In other words, each government department has a set of data on each citizen who is relevant to that department (often the majority of citizens, e.g. department of Health). As each government department is politically distinct, so are the information systems. There is no inherent inter-operability between these departmental silos. Each silo replicates citizen data and has its own physical network and infrastructure. This is not **cost-effective**. We can see that this is true in terms of the replication of hardware and network running costs, but it is also true in the way that each department will have its own bespoke software systems. Not only will the procurement costs be multiplied, but the management costs are also likely to be greater [Max2009].

Some people may argue that these departmental separations are necessary, because the same data has different significance and needs different governance, in each context. However, the Information Cloud gives us intelligent ways to overcome these issues. For example, in the current Internet, it is expected that, to provide a holistic view of a citizen, there would be contention as to which identifier to use to refer to that individual. In the UK, the choice could be between the National Identity Register Number, or the National Insurance Number, or the National Programme for IT (NHS) number. However, the Information Cloud enables the way that information is identified to be flexible, according to context. The Information Cloud enables information (here the set of data associated with an individual) to be referenced in a number of ways, according to what is most appropriate for the end-user. This is because the translation of subscription to publication of data will be handled by the Rendezvous Point, in a way that enables the reference to that information to be expressed within the ontological framework of the information subscriber. With respect to Government, this will remove one of the major motivations for having separate software systems. If we can also remove the need for separate networks and hardware, we will enable a huge step forward in terms of simplification, cost-reduction and inter-operability.

With the PSIRP architecture this is achieved in a number of ways. The Information Cloud decouples information from infrastructure and location. The Information Cloud creates a highly distributed cloud of information items: large and small, which are decoupled from infrastructure and information location. This means that applications can be written purely in information-centric terms. An information-centric network enables access from anywhere through anything, provided access is granted. It makes it easy to write applications that focus on the information: you can re-use the same data for different purposes; and applications are independent from the location of producers and consumers of that information. The second key aspect of the Information Cloud is the way that information is tagged. Tagging provides meta-data on how the information is to be used, for example information source, access rights, relevance, time to live etc. This tagging occurs at multiple levels in the architecture, occurring as 1) secure distribution of keys to give access to RIds, and SIds, where the secure relationship is managed by the Rendezvous point and the RId and/or SId can incorporate policies on information use, which will be read by all nodes in the network as well as end-users ; 2) secure channel for content streaming, whereby indirect cryptography requires the subscriber to trust the publishing host who can implement information policies; 3) metadata expressed in the TTP header of a PLA packet, granting access to authenticated subscribers only. The third most important aspect of the Information Cloud for Government is the way that a new application is just a new set of relationships between existing pieces of information. The network itself and the data remain unchanged. This fundamental difference has far-reaching consequences, and goes together with the distribution of information, to break down departmental boundaries. This also provides enormous benefits in that having only one set of core data means that **consistency** and **accuracy** can be much more easily maintained.

As well as each current government silo having its own citizen records and software, it also has its own security systems. In spite of this security, there have been several scandals in recent years, in the UK, associated with large volumes of data being downloaded on to media and subsequently mislaid, either in the post or in taxis etc. [Gar2009]. This seems to indicate that government employees have access to large amounts of data, but find it difficult to access the specific data that they need using the systems available to them. This seems to be true even within a government department, but the situation is expected to be worse when information is sought across government departments. It is likely that the lack of inter-operability means that to find information on an individual, in a number of contexts, large portions of the database must be downloaded for subsequent analysis off-line. At the very least, the scandals associated with data losses show that there are insufficient constraints on the way that data is accessed, allowing an individual to have access to large volumes of citizen data, when this is probably not strictly necessary, thus leaving the way open for them to use that data carelessly.

Beyond the question of inter-operability, current information systems in most organisations have a poor audit trail. It is difficult to know who has had access to which pieces of data and what happened to it next. When this information refers to the personal records of citizens, there is a particular need for the information framework to provide **accountability**, including the history of creation and modification of records, as well as who has seen those records and how they have been manipulated. Much or all of this functionality can be provided by using the PSIRP architecture with the implementation of Packet Level Authentication. A modified form of PLA provides not only traceability of publishers, but also requires subscriber transactions to be authenticated. This means that there will be a record of all pieces of information that are accessed.

The ability of the system to authenticate and know who has accessed which pieces of information will have a radical effect on security design. We can draw an analogy with the postal system. In the same way that we would not dream of securing every street when the post is being distributed, because the letters are addressed to specific people, securing transactions removes the need to have a fortress approach to information vaults. The fact that transactions must be authenticated, rather than simply authenticating access to a network, removes another driver for having physically separate departmental systems.

A further benefit of removing silo design is in IT system delivery. A common issue with large Government contracts is around fulfilment accountability. Silos make it difficult to monitor cost, to predict demand and to identify inter-dependencies. This is likely to be a big contributor to the way that the costs of IT projects often spiral out of control.

However, the most profound issue with departmental-based information silos is that this system makes it very difficult for anyone to have a holistic view of a citizen. That is, a person's behaviour in one sphere is not easily connected to his/her behaviour in another sphere, so inconsistent assumptions can be made, leaving open opportunities for frauds of various kinds. Here we describe probably the most significant change that the Information Cloud could bring about: **citizen-centricity**.

Current government IT systems are not designed around the needs of the citizen. There is a need for a much stronger user control and ownership framework. In general, the Information Cloud enables a balance of power between producer and consumer of information. There can be equal power by matching availability to interest. In the case of citizen data, the Information Cloud provides the opportunity to give the citizen much more visibility of the data held on him/her. With a citizen-centric information system, one access channel could enable the citizen to view all data held on him/her, ranging from car registration, through tax credits, to medical records. It also opens the way for citizens to be able to update their own data directly, for example when they change their house. Updating their address once would immediately update their address records for all relevant government services.

Here we see a powerful example of the way that rights and responsibilities would be coupled together. As well as the efficiencies enabled by Information Cloud and the likely improvements in consistency and security, we also see a step change in a government's ability to serve its citizens, as well as to prevent fraud. With a holistic citizen-centric system, it will be much more difficult for false identities to be generated, as each citizen will have associated with them a lifetime's set of interactions with government systems, which will be much more difficult to forge. This could make an enormous impact on many branches of national security, as well as reducing the costs of identity fraud.

6.2 Collaborative Business ICT Services

Collaborative Business ICT Services is a group of services enabling businesses to work in more flexible ways, adapting to change at lower risk. This opportunity area has been divided into Retail Multiservice, Dynamic Supply Chains and New Markets, with New Markets expressing the most ambitious and long-term opportunities.

6.2.1 Retail Multiservice

Retail Multiservice refers to a retail business that has diversified beyond being a multi-channel retail organisation, to offering its customers a very diverse range of services. An example of this type of business is TESCO [Cor2009, IVI2009], which, in addition to its grocery business, offers insurance, banking, travel and clothing to its customers. Typically what happens when an organisation wants to diversify into a new type of business, is that it buys up an existing business and re-brands it as their own [MAR2009]. This has the advantage that know-how and business-specific systems are already in place. However, it has the disadvantage that the choice of business, in particular its size, customer base, location etc will be heavily constrained. In general, only the largest companies will be able to carry out this kind of acquisition. Even when the right business has been targeted and acquired, in offering the new business to its existing customer base, customer records and accounts for those existing customers will need to replicated on to the new system. This has implications for time and effort in migration, as well as infrastructure and storage costs. If the acquired business turns out to be unattractive to existing customers, there will be surplus capacity. The risks are high.

What the Information Cloud offers is the ability to reduce the risks associated with diversification, by enabling incremental business growth. In the case of multiservice businesses, this means having a business system which is built entirely with reference to the customer. It is then easy to refer to a customer in any number of different service contexts. A new service is simply a new application operating on the same customer information. Information-centric applications can be written that decouple the location of information from its meaning. There is fully distributed data storage. These factors will result in the ability of a business to enjoy incremental growth of infrastructure and network services and diversification at low cost.

Evidently this means that customer data from different contexts would need to be used in an appropriately secure and reliable way. However, this is something that the PSIRP architecture also gives us, with secure accountable transactions being provided by a modified form of Packet Level Authentication (see section 4.6).

In summary, with the Information Cloud, the same customer data can be tagged in different ways, depending on the service. Associated encryption and policies make it easy to control who can view specific data. This results in something that is revolutionary: we remove the need for service-specific data silos. So the Information Cloud provides the flexibility to handle new services and customer growth, without significant risk to the business.

6.2.2 Dynamic Supply Chains

Supply chains work well when there is predictable demand and reliable suppliers to meet that demand. In reality, businesses need to constantly update the range of products that they offer; some suppliers may not be able to meet demand, or may be otherwise unreliable.

A common issue revolves around how to trial new product ranges. It would be highly advantageous to a business if it was easy and low-cost to interface new suppliers with corporate operational systems to enable short trial periods.

Another important aspect for a business is how to cope with supplier failure to deliver, whether this is due to cash flow problems, or raw material supply problems, or local weather damage or logistical problems etc. An ideal situation is for dynamic supply chains to enable adaptation to these short-term or longer-term difficulties. With the Information Cloud, it is easy to configure time-limited and policy-constrained virtual relationships amongst suppliers. This means that small suppliers can join a business for a short period in a loosely-coupled way. In fact it means that all supplier relationships are loosely coupled, with interfaces that capture the key state associated with that part of the supply chain, such as current and predicted stock levels and also indicators of stability related to cash flow and operational issues etc. The supplier relationships are then easy to set up, as opposed to the current state of affairs where a supplier has to effectively join the retailer's corporate system. In addition, it is envisaged that the supplier interfaces, which capture a supplier's ability to meet its commitments, will be made available, in a controlled way, to other potential suppliers, up and down the supply chain. This will enable self-organising dynamic supply-chains, whereby suppliers at various levels of the chain can step in to back-fill gaps in the supply chain, before or as they occur. This should mean that the balance between supply and demand, at every level of the supply chain is as well-matched as possible. Such an ecosystem would have to be regulated, of course, but would not need the slow, manual contract-negotiation and heavy system set-up costs of current arrangements. The efficiency associated with being able to anticipate and fill gaps in the supply chain and the consequent ability to serve demand, will give companies participating in dynamic supply chains a significant edge over their competition.

6.2.3 New Markets

When referring to New Markets, we could mean any business expansion, but in particular we are thinking of significant expansion into a new country or new continent. This is probably the most risky business decision, as it is likely to be large in scale, and will either involve buying a new company, or trying to combine existing business practices and systems with the regulations and market behaviours of the foreign country. In considering this expansion, we need to think about the entire set of end to end business ICT systems, and how they can be simplified.

The Information Cloud has as its premise that everything is information. This helps us to approach the challenge of entering a foreign market by considering that applications, software components, customer data and operational state of all kinds, can be treated as items of information in the cloud. So when we think of operations such as procurement, delivery, billing and stock-taking etc, the goods being sold as well as the services used to sell them are all handled within the same information framework.

When we considered dynamic supply chains, we saw the importance of creating flexible relationships amongst suppliers. We now take this a stage further and consider how to build flexible relationships amongst pieces of operational code. This would mean that old services in new markets could be made operational very rapidly, as we describe in more detail below.

Treating software components as items in the information cloud means that software components are decoupled from infrastructure and location. Again, following the analogy of peer-to-peer networks, this inherent distribution and mobility paves the way for software components to proliferate and migrate towards user demand.

In moving into a foreign market, new components will be needed, to meet new regulatory and market pressures. However, writing services as software components means that many component elements can be re-used. With the Information Cloud, policies can be expressed to constrain the choice of component elements and the way that they are put together, to optimise service operations. We can think of these components as web services. Their properties could be their reliability, bandwidth, latency or cost etc. This should enable smooth service operation when moving into new network infrastructures. It might even enable self-organising services.

With the Information Cloud, transaction accounting is easy because it is an inherent part of the way that the network operates. This applies to both customer transactions and system transactions, that is, both the purchase of goods and the use of operational applications to service those requests. For example, with the Information Cloud, the access to information is inherently audited. This could revolutionise the way that stock-taking takes place, by linking purchases right back through the supply chain. More generally, the auditing of access to applications means that flows of service attributes, the performance of service components, their usage and billing can be captured without needing additional system layers. As well as this being useful for managing stock and cash-flow levels, it is important at a more profound level. Understanding the pressure points of a system is critical for successful sizing of operational systems. Operating within an Information Cloud gives a business an in-depth understanding of how its business systems are being used and so enables rapid adaptation.

In summary, applying the information cloud approach to software, as well as data, means that business systems comprising software components, system storage and ultimately the network capacity itself can adapt and potentially self-organise in response to changing demand.

The ability for business systems, at every level, to adapt to market pressure means that spare capacity can be kept to a minimum and costs can grow with market share. This gives a huge cost advantage in dimensioning operational systems in new markets and massively reduces business risk.

6.3 Content-Centric

We have seen in the previous application opportunity areas, that a key aspect of the Information Cloud is in the way that it stores information, by forming a highly distributed cloud of information items. There is decoupling of information from infrastructure and location, which means that applications can be written in information-centric terms. Secondly, the fact that information is organised, accessed and routed according to its meaning enables information to migrate towards demand. This means that application performance can be optimised and the use of infrastructure can be optimised. Thirdly, the Information Cloud enables a rich set of tags to be used to control the way that information is handled in the network and how it can be used by subscribers.

What these three main characteristics enable together is flexibility. The Information Cloud enables new services and applications to be created purely by reconfiguring the elements within the cloud. In referring to an element in the cloud, we understand that this is not just a piece of data, although it could be, it can also be a much bigger collection of bits, such as a media file; and it can also be a piece of code. Along with the flexibility of elements that can be configured together, there is a richness that is achieved by expressing policies about the ways in which the elements can be configured together, for how long and by whom. This results in the ability of the Information Cloud to respond very rapidly to changing requirements, which is key to enabling new responsive media experiences, as we will see in the examples that follow.

For new applications in a media context, we consider that major drivers are Empowerment, Truth and New Realities. We take as inspiration for this analysis, other work on the Future Content-Centric Internet [COR2009, FCN2009].

6.3.1 Empowerment

During the twentieth century, important events happening to the general public were typically captured by taking photographs. With the advent of the Internet, sites such as YouTube have flourished, giving people the ability to also share film footage with each other. Film footage of world events forms an important part of this. Amateur footage taken when dramatic events take place and there are no film crews around becomes highly valuable. The quality of footage is more than the number of pixels. There is the desire for community-based content and for quirky/artistic/humorous footage taken from odd perspectives and people do wish to use multimedia to participate in important events. However, at some mass events, such as the Olympics, it will not be desirable for thousands of people in a stadium to simultaneously upload their footage of gold-medal winning performances to YouTube or other sites, as much of this film would be near-identical and of low quality, each with a very limited audience, therefore the use of bandwidth for this is questionable.

So, can the Internet operate in such a way as to match availability to audience? We think it can, if it operates by using an Information Cloud (PSIRP) architecture. This is because the PSIRP architecture operates in a publish/subscribe way, at multiple levels. This means, in essence, that content without an audience will not take up bandwidth, because it has no subscribers. Content that is in demand will proliferate and be easy to access. Content with few subscribers will be inherently less well served. That does not mean that the content cannot be published, it just means that it will consume fewer resources, as its publication will remain local to where the content was generated.

Another major aspect around ‘empowerment’ with respect to content is the ease with which content can be discovered. At the moment, it is often important to know the name of a website, in order to find the content you want. Without this, an individual is left trying to guess the best search-terms to hit on the content they want. The Information Cloud should help to make this significantly easier.

The Information Cloud will enable anyone to set up a Rendezvous Point (RP). This RP will then act as the gateway to content which has in common a certain reputation or subject-matter or political perspective etc.; or, in other words, falls within a certain scope. Remember that pieces of content may be in multiple scopes. Whilst running a Rendezvous Point may seem the same as hosting a website, with links to other websites, there are a number of key differences. Firstly, there is the difference of information being decoupled from location. This means that the structure of the Rendezvous Point can be much more flexibly configured than a website, and the connections that it makes are not tied to physical locations, either, so they will be more reliable. Secondly, the hierarchy of scopes, where all information is organised according to its meaning, should result in groupings of information according to their meaning, as viewed by users of that content.

Rather than requiring hands-on editorship, the membership of scopes could actually be allowed to self-organise, whereby content that falls out of scope will naturally be sidelined, in favour of the content that people going to that RP have subscribed to view. For some Rendezvous Points, the risk of unsuitable material being accessed will be unacceptable and moderating the content will be imperative. However, that is also true for a website. The difference with a Rendezvous Point is that a reputation for good moderation, good governance and high visibility etc will draw publishers to the Rendezvous Point. This is in contrast to the current Internet, where links are put in place by the host website by hand, and content would need to be forwarded and individually proposed to the website owner. With the Information Cloud, opportunities would exist for mutual trust relationships to develop through networks of rendezvous points, such that reputable providers of content become associated together. With this greater fluidity of associations of content, it is envisaged that not only will content be easier to locate, but the most ethical or most effective rendezvous points will flourish, as a true reflection of their effectiveness and/or values. Or, at least, they will be associated together so that subscribers know whom they can trust to demonstrate a certain quality of content provision.

A final element of empowerment for a content-centric Internet is around censorship of content. There are situations, for example, the classroom, where it is desirable to limit access to certain unsuitable content. At the moment, the types of filters that can be applied are fairly crude, according to general subject-matter headings. This can result in very useful content being barred, because it cannot be distinguished according to these high-level search terms. With the Information Cloud, a scope could embody an ethical perspective, or at the least will be capable of applying much more sophisticated policies. Whilst the onus could be argued to be on the subscriber to vet the access to content, a Rendezvous Point would enable content to be tagged as being unsuitable for certain audiences, thereby making it much easier for educators to know what is safe to show in the classroom, without blocking an entire website. This capability can be seen in the context of information security, where it is the information which is secured, not the physical network or website.

A closely-related set of characteristics that will drive the new Internet architecture can be summarised as being related to the authenticity of content, or ‘truth’, as is described in the next section.

6.3.2 Truth

Currently, amateur news footage typically reaches its audience via a news agency. The audience relies on the news agency to be reputable and to carry out due diligence to establish the authenticity of that footage. This approach has stood the test of time, and may be the best model. However, it is also a quite limited approach. What if access to international news agencies is blocked by a repressive regime, as is access to specific sites, such as YouTube? The Information Cloud opens the way for many more publication opportunities, which would be much harder to control.

The Information Cloud would also enable people with less momentous stories to make an impact, via publication in smaller interest groupings of content. This would be particularly relevant for whistle-blowers. A whistle-blower is someone who observes some kind of wrongdoing and who does not have confidence in the immediate organisation to take the appropriate action. They may feel themselves to be in danger of losing their job or being otherwise penalised if they speak out. So they need a way to bring the information to the attention of outsiders, who have an impartial view, without jeopardising their own position. Publication to a Rendezvous Point that specialises in, for example, hospital hygiene issues, could enable them to bring to wider notice evidence of poor hospital practice. In this case, the RP would know the identity of the originator of the information, so that authorities, such as the police, could subsequently investigate the validity of the allegations. However, the whistle-blower’s identity would be protected from the attention of the employer.

Authenticating digital film or other media is difficult. The Information Cloud may not have all the answers, but it does offer significant benefits. Rights management is of critical importance to the creative content industry. The Information Cloud enables a much tighter coupling between content and the consumers who have paid to view that content. The way that content is accessed would be much more tightly controlled. Authentication of the subscriber would be needed, as well as verification that the content is original and not pirated. It would mean that we can effectively licence each piece of media to the person who bought it and thereby prevent the content from being passed on to or shared by others. This has huge implications for the creative content industry, which struggles with the practical and legal steps needed to protect access to content. We believe, therefore that it is an important driver for the adoption of PSIRP and PLA approaches.

The process of authenticating content with the Information Cloud would also enable much greater control over content manipulation. Key images within a frame could be separately annotated, such as people's faces, to prevent them from being 'photoshopped'. This would help to prevent people from being represented carrying out actions that were carried out by others. The current ease of image manipulation and its lack of traceability may explain why multimedia are mostly absent from sites such as Wikipedia. The ability of the Information Cloud to tag subsets of images, for example faces, would prevent those images from being republished out of context. This applies not only to photoshopping but also to the juxtaposition of images. For example, religious groups would be unhappy to have religious images displayed alongside blasphemous images. The Information Cloud would enable publishers to know that their content has been incorporated into a scope that they do not approve of, and to remove it from that scope.

The final application area that we consider within content-centric opportunities can be described as enabling New Realities. Here we consider a few examples and their requirements.

6.3.3 New Realities

In summary, we believe that the Information Cloud will make real-time 3-D media mash-ups much easier to create. Of course, the network itself cannot provide everything to achieve these applications. Other technologies would also be needed, such as audio-visual analysis techniques and component preserving capture technology. However, the network has an important role to play.

We discussed earlier that the way that information is tagged will enable meta-data to be provided to end-users and their devices. This will be key to rendering 3-D models created by user actions. For immersive environments to be successful, another critical issue is network performance. The Information Cloud means that gaming software components and player data can move towards the players to optimise the gaming experience. Another key factor, which applies to all New Reality opportunities, is the ability to reserve high bandwidth links for the duration of that reality. The Information Cloud enables information to be tagged, via the RID, that could be used to indicate that certain traffic is to be given the highest priority. The fact that this instruction is encoded in the packet header means that there is no need for routers to carry out deep-packet inspection or to have knowledge of the subscribers. It is enough that they have a relationship with the publisher, who manages revenue gathering from the subscribers, who are given the preferential interaction experience. An obvious application of this technology is for interactive 3-D gaming.

Other applications that would rely on high-bandwidth connections and secure content are haptic interactions, for example for the teaching of high-level skills such as veterinary studies. Here, modelling of 3D space and user interactions will need to be conveyed with high fidelity, fast and reliably.

A similar set of requirements exists for co-operative new realities, such as augmented reality for product design. Here virtual teams will need almost instant visual and tactile feedback from product components that they can touch and view simultaneously.

We will now review the trigger points to enable Information Cloud to deliver on these Content-Centric visions.

7. Business Modelling

So far we have captured stories that illustrate business opportunities in three sectors. However, we would like to be able to quantify how these business opportunities will interplay and lead to the transformation of the network, given certain assumptions of cost and benefit. Therefore, the key opportunity areas have been used to develop a business model, primarily from the perspective of an incumbent telecommunications operator. Real financial data has not been used to parameterize the model, so there is nothing specific to an operator or to a market. In developing the model, we are chiefly interested in relative costs and benefits, and what are the tipping points to enable the PSIRP architecture to be taken up, according to evolving business benefits. If the key opportunity areas are transferable to other markets, then so will be the business model.

The model operates over up to 30 years, starting from a point where there is some demand for the key opportunity areas, but as yet no capacity for the network to fulfil those requirements. Each key opportunity area depends on certain technology enablers. Below we identify what we see as the most important technology enablers to the key opportunity areas, many of which will be modelled within the systems dynamics approach and so are referred to as trigger points.

7.1 Government Trigger Points

- Government Rendezvous Points
    - Hierarchy of scopes (managed by Rendezvous Points) to map to government departments
    - Development of appropriate ontology handling so that employees can access information in the right context easily
- Government owned Information Cloud infrastructure, within which information is distributed
- Auditing of Information Transactions
- Packet-Level Authentication to enable accountability of data access and to ensure information is authentic
- Algorithmic RIds for versioning of documents in caches
- Ultimately, secure channels for RID distribution to enable citizens to access and update their own records. This could be mediated via interactions with the NHS.

7.2 Collaborative Business ICT Service Trigger Points

- Corporate virtual private information cloud, via shim layer
- Corporate Rendezvous Points
- Packet Level Authentication to enable secure and accountable access to customer information in a variety of business contexts
- Algorithmic identifiers to maintain version control in information caches
- Ability to write operational code in appropriate service components.
- Mechanisms to enable code proliferation and migration
- Mechanisms to enable service composition
- Auditing of software and network usage

- Auditing of information transactions

7.3 Content-Centric Trigger Points

- Virtual private information clouds, via shim layer initially
  - Educators
  - News agencies
  - Research institutes
- Virtual private information clouds for mass markets, via shim layer initially
  - Content producers and distributors
- Rendezvous Points
  - Broadcasters
  - Content producers and distributors
- Secure channels for RID distribution
  - Content producers and distributors and retailers
- Indirect cryptography to handle content distribution channels
- Packet Level Authentication to enable authentication of subscribers and content
- Native PSIRP solutions with widespread router implementation of RID reading for packet-level switching for high bandwidth applications
- Algorithmic identifiers to maintain version control in information caches
- Mechanisms to enable code proliferation and migration
- Mechanisms to enable service composition
- Audio-visual analysis and component-preserving capture technology

Most of the technology enablers are capabilities made possible by the PSIRP architecture. The most important of these trigger points (sometimes expressed as composites of those listed above) have been modelled. The modelling of the trigger points involves their requiring investment, on the one hand, and generating demand for traffic, on the other.

7.4 Deployment Strategies

The model is designed to show the evolution of solutions to demand from application sectors. Adoption of the Information Cloud has been separated into two solutions: the first is referred to as a ‘Shim Layer Information Cloud’. This is expected to be a tightly coupled overlay approach, where IP is still the underlying network protocol, but information can be accessed and handled using information identifiers. This is expected to be the initial and interim solution. The second solution we refer to as ‘Native Information Cloud’ in which the nodes in the network are operating fully as PSIRP nodes, with IP no longer being necessary. We do, however, envisage both these solutions co-existing with IP for many years, so this is enabled in the model.

It will be seen that most trigger points are modelled as generating demand for ‘shim layer’ traffic, because their service needs can be met by this type of solution. However, we envisage that a ‘native’ PSIRP solution, whilst requiring more infrastructure investment, will ultimately provide higher revenue than ‘shim layer’ solutions, so we also model the migration of traffic from ‘shim’ to ‘native’.

The migration from one type of network to another is the migration that is most relevant to the telecommunications operator. However, this is driven by the functional migration from isolated services, to services within an industry, leading, finally, to what can be classed as an ‘internet of services’. An internet of services is what a native PSIRP architecture fundamentally provides, which cannot be delivered using IP or IP overlays. This is represented in the model as the trigger point drivers ‘dynamic service composition’ and ‘dynamic bandwidth reservation by SId/RId’ which, in our model, cannot be delivered by a ‘shim layer’ approach and therefore feeds into the demand for a native information cloud only. In principle, dynamic bandwidth reservation could be achieved using MPLS and deep packet inspection, but in practice this would be impractical for all but the highest value applications.

7.5 Outline of Model

By connecting investment in key opportunity areas to their technology trigger points; and modelling the generation of different types of traffic, ISP profit and hence infrastructure investment, we hope to be able to model the critical factors in the life cycle of PSIRP architecture adoption.

First of all we present the basic set of relationships in the model, as a Vensim system dynamics model, shown in Figure 7.1.

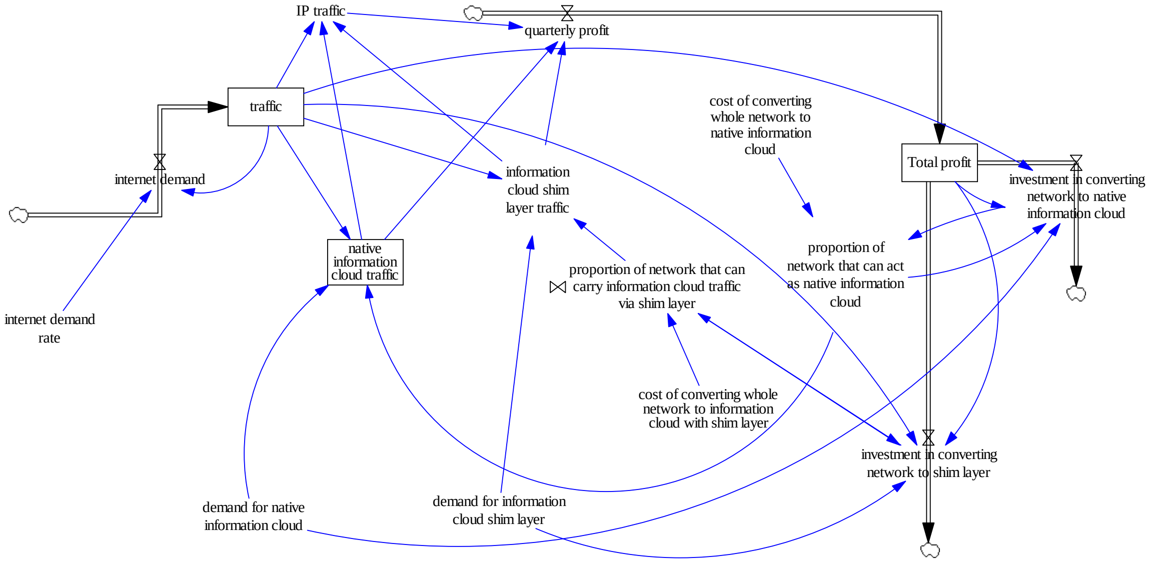
IP traffic
quarterly profit
traffic
cost of converting whole network to native information cloud
Total profit
internet demand
information cloud shim layer traffic
investment in converting network to native information cloud
native information cloud traffic
proportion of network that can carry information cloud traffic via shim layer
proportion of network that can act as native information cloud
internet demand rate
cost of converting whole network to information cloud with shim layer
investment in converting network to shim layer
demand for native information cloud
demand for information cloud shim layer

Figure 7.1: Outline of system dynamics business model

We see that there is a positive feedback loop associated with Internet traffic generally. This traffic can either be IP, ‘information cloud shim layer’ or ‘native information cloud’. All types of traffic generate profits for the Internet Service Provider/Telecom operator, on a quarterly basis. Total profit is a stock whose level rises with quarterly profit, but out of which investments are made. The investments we model are to convert the network to be able to ‘carry information cloud traffic via shim layer’ and separately to enable the network to ‘act as native information cloud’. It is expected that implementing native information cloud will be more expensive than the shim layer, but that native information cloud traffic will be the most profitable, compared with shim layer, and IP traffic.

In this initial simplified system, we do not see what causes the demand for the new types of traffic. We see investment creating capacity and changing demand which is then able to make use of that capacity, to generate traffic.

7.6 Business Model in Detail

In the more complete model, we are able to identify investment from the key opportunity areas driving demand for Information Cloud traffic, via specific functionalities, which are critical to the delivery of those applications.

In Figure 7.2 we see the key contributors to ‘native information cloud traffic’ within the model. There are three trigger points that directly generate native traffic: “dynamic bandwidth reservation by SID/RID”, “management of new democratic TV formats” and “dynamic service composition”, all of which we feel will directly affect the native information cloud traffic levels, as these technologies cannot be served by an IP overlay or ‘shim layer’. In the model, the trigger point “incremental bandwidth and infrastructure growth” is heavily dominated by native information cloud traffic levels, with a much smaller dependence on shim layer traffic. However, although “incremental bandwidth and infrastructure growth”, is strongly dependent on native information cloud traffic levels, it only has a ‘feed-forward’ effect on shim layer traffic. This choice was made because, although the native information cloud network is very important in realizing the functionality “incremental bandwidth and infrastructure growth”, this technology is not expected, in itself, to generate high volumes of traffic.

Note that as the network is transformed so as to be able to carry native information cloud traffic, shim layer traffic is assumed to be converted to native traffic. This means that the drivers for shim layer traffic effectively become drivers for native traffic, as the network is transformed. The justification for this, in modelling terms, is twofold. First, we want to be able to distinguish between technologies that depend on native Information Cloud solutions, and those that can be served by an interim ‘shim layer’ solution. However, secondly, we want to be able to show that the native information cloud can provide everything that the shim layer information cloud can provide (and more) so that, if there is an incentive for the network operator to provide native solutions, in terms of greater revenues, then it is assumed that shim layer solutions will gradually be phased out and replaced with native solutions. For simplicity of network management, and in order to offer the greatest functionality to customers, having a network that is pure native Information Cloud seems advantageous. However, we acknowledge that this is an assumption of the model, and that if relative costs and benefits of the two solutions have been inaccurately modelled, then the conversion of shim layer solutions to native solutions may be much more limited.

The volume of native information cloud traffic is a key measure of success of the migration, in our model. There are many contributors to its magnitude, as illustrated in the tree diagram below.

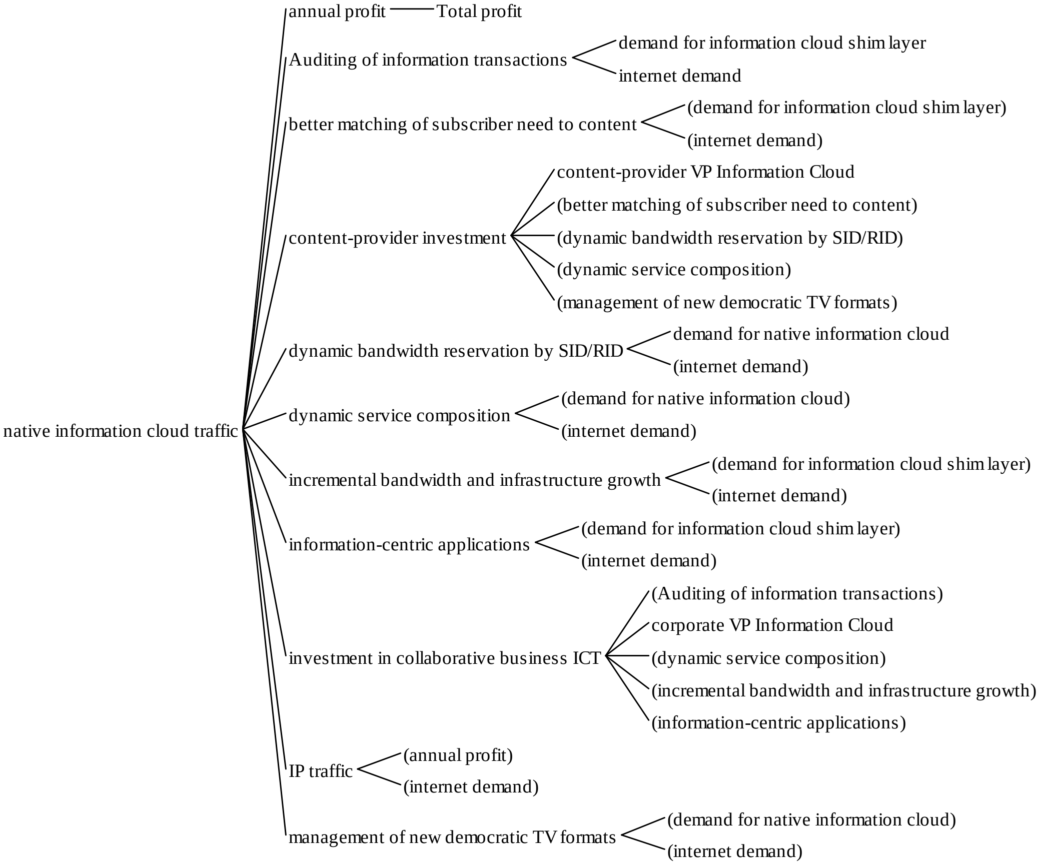
native information cloud traffic
annual profit
Total profit
Auditing of information transactions
demand for information cloud shim layer
internet demand
better matching of subscriber need to content
(demand for information cloud shim layer)
(internet demand)
content-provider investment
content-provider VP Information Cloud
(better matching of subscriber need to content)
(dynamic bandwidth reservation by SID/RID)
(dynamic service composition)
(management of new democratic TV formats)
dynamic bandwidth reservation by SID/RID
demand for native information cloud
(internet demand)
dynamic service composition
(demand for native information cloud)
(internet demand)
incremental bandwidth and infrastructure growth
(demand for information cloud shim layer)
(internet demand)
information-centric applications
(demand for information cloud shim layer)
(internet demand)
investment in collaborative business ICT
(Auditing of information transactions)
corporate VP Information Cloud
(dynamic service composition)
(incremental bandwidth and infrastructure growth)
(information-centric applications)
IP traffic
(annual profit)
(internet demand)
management of new democratic TV formats
(demand for native information cloud)
(internet demand)

Figure 7.2: Tree diagram of contributing factors to the level of 'native' Information Cloud traffic

In our model, we consider the investment needed to transform the network, but also the investment needed from the key application sectors. As described previously, these are Government, Collaborative ICT and Content-Centric. These application investments are linked to technology triggers and traffic levels as follows:

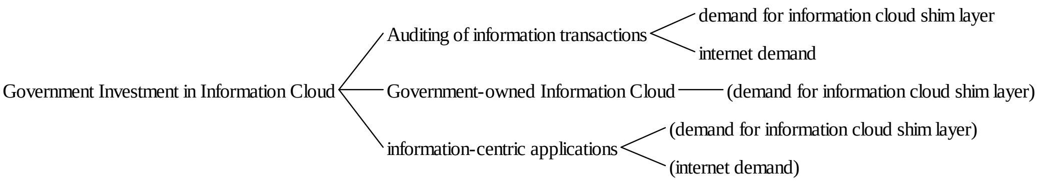
Government Investment in Information Cloud
Auditing of information transactions
demand for information cloud shim layer
internet demand
Government-owned Information Cloud
(demand for information cloud shim layer)
information-centric applications
(demand for information cloud shim layer)
(internet demand)

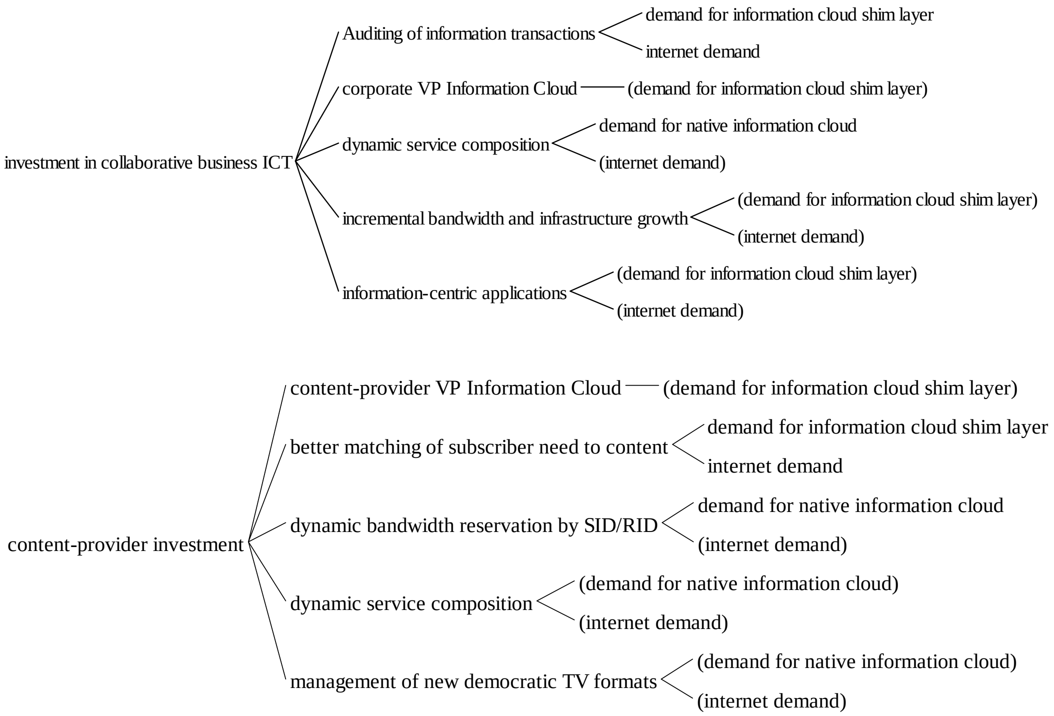
investment in collaborative business ICT
Auditing of information transactions
demand for information cloud shim layer
internet demand
corporate VP Information Cloud
(demand for information cloud shim layer)
dynamic service composition
demand for native information cloud
(internet demand)
incremental bandwidth and infrastructure growth
(demand for information cloud shim layer)
(internet demand)
information-centric applications
(demand for information cloud shim layer)
(internet demand)
content-provider investment
content-provider VP Information Cloud
(demand for information cloud shim layer)
better matching of subscriber need to content
demand for information cloud shim layer
internet demand
dynamic bandwidth reservation by SID/RID
demand for native information cloud
(internet demand)
dynamic service composition
(demand for native information cloud)
(internet demand)
management of new democratic TV formats
(demand for native information cloud)
(internet demand)

Figure 7.3: Contributing factors to application sector investments

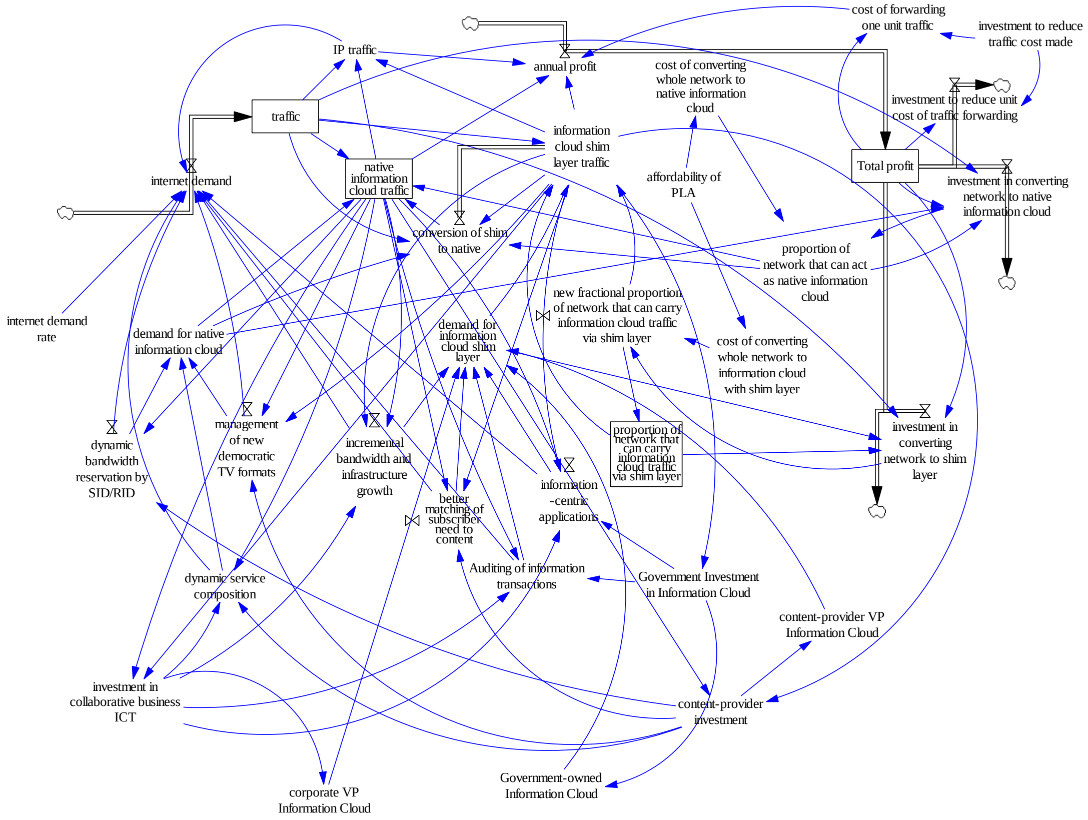
cost of forwarding one unit traffic
investment to reduce traffic cost made
IP traffic
annual profit
cost of converting whole network to native information cloud
investment to reduce unit cost of traffic forwarding
traffic
information cloud shim layer traffic
Total profit
internet demand
native information cloud traffic
affordability of PLA
investment in converting network to native information cloud
conversion of shim to native
proportion of network that can act as native information cloud
internet demand rate
demand for native information cloud
new fractional proportion of network that can carry information cloud traffic via shim layer
demand for information cloud shim layer
cost of converting whole network to information cloud with shim layer
management of new democratic TV formats
dynamic bandwidth reservation by SID/RID
incremental bandwidth and infrastructure growth
proportion of network that can carry information cloud traffic via shim layer
investment in converting network to shim layer
information -centric applications
better matching of subscriber need to content
Auditing of information transactions
Government Investment in Information Cloud
dynamic service composition
content-provider VP Information Cloud
investment in collaborative business ICT
content-provider investment
Government-owned Information Cloud
corporate VP Information Cloud

Figure 7.4: System dynamics business model in detail

We see that the Government requirements can be met from a shim layer solution, but both the other two application areas require a native information cloud to be available in order for all their technology needs to be met.

In Figure 7.4 we see a much fuller description of the model, in which many of the trigger points that were identified whilst looking at the key opportunity areas are modelled. We explicitly enter the cost of providing packet level authentication (PLA) as a parameter that needs to be applied to either Information Cloud solution, as we believe that it is key to the vast majority of application opportunities using PSIRP. We also provide for shim layer traffic to be converted, over time, into native information cloud traffic, as this will maximize profitability for the infrastructure owners. We model a lower cost of network transformation to shim layer, compared with the native solution, but also a lower level of profit per unit of traffic flow. IP traffic is the default type of traffic, which, without network investment, will grow, but it has the lowest level of return for the infrastructure owner. What we generally see is an initial predominant conversion of the pure IP network to a shim layer, a period in which native information cloud traffic begins to grow, and an ultimate reduction in shim layer traffic, as native traffic takes over from the shim layer.

7.7 Scenario

The principle scenario being modelled is one in which three major sectors are modelled as providing the demand for traffic and new services over a PSIRP architecture. We assume that the potential investment that the three sectors provide is conceptually independent. However, there is also a strong feed-forward component in investment, which relies on network traffic levels, which can be stimulated by other players.

The investments from the key sector areas: Government, Content-Centric and Collaborative ICT, each have a constant component and a component that depends on traffic volumes. Note also that, in each case, annual investment is ultimately capped, despite exponentially rising traffic levels.

We assume that the network begins with IP traffic only, generating a certain amount of revenue and profit for the network operator. We assume that demand for network traffic grows yearly, but that this rate increases when new functionalities are enabled, as a result of PSIRP technology.

We assume that there will be regulation of profit, so that profit levels rise much more slowly than traffic volumes and will sometimes be arbitrarily capped. Traffic costs rise linearly, and so will tend to overtake traffic revenues. However, we also make provision for the network operator to undertake further investments to reduce the unit cost of forwarding packets. In this way, we try to be realistic about the potential for increased profits from new products and services, whilst modelling the very large rises in traffic volumes enabled by new capabilities.

We map the sector requirements to demand for specific technology components of functionality that the PSIRP architecture could deliver. We distinguish between technologies that rely on a ‘native’ PSIRP architecture, and those that could be provided by a ‘shim layer’ network overlay, or – in other words – to a large extent by off the shelf components available to network operators today. As previously mentioned, there is a feed- forward effect in driving demand for traffic, according to investment in providing specific new functionalities enabled by the information-centric networking and publish-subscribe paradigm that PSIRP encompasses. Note that there is a slightly different treatment of investments in VP or Virtual Private networks, in which solutions are provided for a specific sector only; in this case, the feed-forward effect is lower, as this is not expected to stimulate network traffic in other sectors.

7.7.1 Definition of Model Parameters

Below is a list of equations for all the variables in the model:

(01) affordability of PLA=

15

Units: **undefined**

(02) annual profit=

SQRT(SQRT(SQRT(SQRT(SQRT(SQRT(IP traffic*0.04))))))+IF THEN ELSE((information cloud shim layer traffic

+native information cloud traffic

)<1e+006, (SQRT((information cloud shim layer traffic*0.002)+(native information cloud traffic

*0.003))) , IF THEN ELSE(

(information cloud shim layer traffic+native information cloud traffic)<1e+008

,SQRT((information cloud shim layer traffic

*0.0004)+(native information cloud traffic*0.0006)) , SQRT((information cloud shim layer traffic

*0.0004)+(native information cloud traffic

*0.0006))))- ((information cloud shim layer traffic+IP traffic+native information cloud traffic

)*cost of forwarding one unit of traffic)

Units: **undefined**

(03) Auditing of information transactions=

(0.4*Government Investment in Information Cloud)+0.1*(investment in collaborative business ICT

)+(0.04*native information cloud traffic)

Units: **undefined**

(04) better matching of subscriber need to content=

0.08*information cloud shim layer traffic+(0.2*"content-provider investment"

)+(0.08*native information cloud traffic)

Units: **undefined**

(05) "content-provider investment"=

IF THEN ELSE(native information cloud traffic<1e+006, 100+(0.01*information cloud shim layer traffic

)+(0.015*native information cloud traffic

),150000)

Units: **undefined**

(06) "content-provider VP Information Cloud"=

INTEG(0.1*"content-provider investment",0)

Units: **undefined**

(07) conversion of shim to native=

IF THEN ELSE(demand for native information cloud<(proportion of network that can act as native information cloud

*traffic), (proportion of network that can act as native information cloud

*traffic)-demand for native information cloud

, 0)

Units: **undefined**

(08) corporate VP Information Cloud=

0.08*investment in collaborative business ICT

Units: **undefined**

(09) cost of converting whole network to information cloud with shim layer=

affordability of PLA+20

Units: **undefined**

(10) cost of converting whole network to native information cloud=

affordability of PLA+30

Units: **undefined**

(11) cost of forwarding one unit of traffic=

IF THEN ELSE((Total profit > 1000), 8e-008,IF THEN ELSE((Total profit>300

), 2e-007 , IF THEN ELSE(Total profit >200, 4e-007, 2e-006)))

Units: **undefined**

(12) demand for information cloud shim layer= INTEG (

better matching of subscriber need to content+incremental bandwidth and infrastructure growth

+"information-centric applications"+Auditing of information transactions+"content-provider VP Information Cloud"

+corporate VP Information Cloud+"Government-owned Information Cloud",

16)

Units: **undefined**

(13) demand for native information cloud= INTEG (

"dynamic bandwidth reservation by SID/RID"+management of new democratic TV formats

+dynamic service composition,

10)

Units: **undefined**

(14) "dynamic bandwidth reservation by SID/RID"=

(0.3*"content-provider investment")+(0.12*native information cloud traffic

)

Units: **undefined**

(15) dynamic service composition=

0.3*"content-provider investment"+(0.4*investment in collaborative business ICT

)+(0.14*native information cloud traffic)

Units: **undefined**

(16) FINAL TIME = 10

Units: Year

The final time for the simulation.

(17) Government Investment in Information Cloud=

IF THEN ELSE(information cloud shim layer traffic<1e+007, 100+(0.005*information cloud shim layer traffic

),55000)

Units: **undefined**

(18) "Government-owned Information Cloud"=

0.1*Government Investment in Information Cloud

Units: **undefined**

(19) incremental bandwidth and infrastructure growth=

0.04*information cloud shim layer traffic+(0.02*investment in collaborative business ICT

)+(0.01*native information cloud traffic)

Units: **undefined**

(20) information cloud shim layer traffic=

IF THEN ELSE(conversion of shim to native >(demand for information cloud shim layer

),0,IF THEN ELSE(demand for information cloud shim layer>(proportion of network that can carry information cloud traffic via shim layer

*traffic) , (proportion of network that can carry information cloud traffic via shim layer

*traffic)-conversion of shim to native, demand for information cloud shim layer

-(conversion of shim to native)))

Units: **undefined**

(21) "information-centric applications"=

0.06*information cloud shim layer traffic+(0.3*Government Investment in Information Cloud

)+(0.1*investment in collaborative business ICT

)+(0.06*native information cloud traffic)

Units: **undefined**

(22) INITIAL TIME = 0

Units: Year

The initial time for the simulation.

(23) internet demand=

(internet demand rate*traffic)+better matching of subscriber need to content

+"dynamic bandwidth reservation by SID/RID"

+incremental bandwidth and infrastructure growth+"information-centric applications"

+management of new democratic TV formats

+dynamic service composition+Auditing of information transactions+"content-provider VP Information Cloud"

+corporate VP Information Cloud+"Government-owned Information Cloud"

Units: **undefined**

(24) internet demand rate=

0.1

Units: **undefined**

(25) investment in collaborative business ICT=

IF THEN ELSE(native information cloud traffic<2e+007,6000+(0.03*information cloud shim layer traffic

)+(0.005*native information cloud traffic

),100000)

Units: **undefined**

(26) investment in converting network to native information cloud=

IF THEN ELSE((Total profit>15):AND:(proportion of network that can act as native information cloud

<0.98

) :AND:((demand for native information cloud/traffic)>(0.7*proportion of network that can act as native information cloud

)), 10, 0)

Units: **undefined**

(27) investment in converting network to shim layer=

IF THEN ELSE((Total profit>10):AND:(proportion of network that can carry information cloud traffic via shim layer

<0.98

) :AND:((demand for information cloud shim layer/traffic)>(0.7*proportion of network that can carry information cloud traffic via shim layer

)), 5, 0)

Units: **undefined**

(28) investment to reduce unit cost of traffic=

IF THEN ELSE((Total profit > 1000),400,IF THEN ELSE((Total profit>300),

150, IF THEN ELSE(Total profit >200, 50, 0)))

Units: **undefined**

(29) IP traffic=

IF THEN ELSE((traffic-information cloud shim layer traffic-native information cloud traffic

>1),traffic-information cloud shim layer traffic

- native information cloud traffic,0)

Units: **undefined**

(30) management of new democratic TV formats=

0.14*native information cloud traffic+(0.1*"content-provider investment")

+(0.06*information cloud shim layer traffic)

Units: **undefined**

(31) native information cloud traffic=

IF THEN ELSE((demand for native information cloud+conversion of shim to native

)>(proportion of network that can act as native information cloud

*traffic) , proportion of network that can act as native information cloud

*traffic , demand for native information cloud

+conversion of shim to native)

Units: **undefined**

(32) proportion of network that can act as native information cloud= INTEG

(

investment in converting network to native information cloud/cost of converting whole network to native information cloud

,

0)

Units: **undefined**

(33) proportion of network that can carry information cloud traffic via shim layer

= INTEG (

investment in converting network to shim layer/cost of converting whole network to information cloud with shim layer

,

0)

Units: **undefined**

(34) SAVEPER =

TIME STEP

Units: Year [0,?]

The frequency with which output is stored.

(35) TIME STEP = 0.25

Units: Year [0,?]

The time step for the simulation.

(36) Total profit= INTEG (

annual profit-investment in converting network to native information cloud

-investment in converting network to shim layer-investment to reduce unit cost of traffic

-investment to reduce unit cost of traffic,

```
                  10)
        Units: **undefined**


(37)    traffic= INTEG (
                internet demand,
                        100)
        Units: **undefined**
```

7.8 Sub-Set Scenarios

Given the scenario as defined in the previous section, a sub-set of scenarios has been developed in which one of the parameters is varied at a time, in order to understand its effect on the ultimate outcome of the model. Remember that the main aim of this is to understand how the investment decisions of the market and the network operator will influence the economic success of that network operator.

The levels of investment provided by the application providers and by the network, are seen as the key overall parameters that govern the success of uptake of the Information Cloud and hence the profitability of the Telecommunications operator. Investment levels for the application providers have been categorized as either ‘proactive’ or ‘sceptical’. A set of scenarios in which investment players adopt different investment levels, is explored. Network investment is the other main variable, where we vary the Total Profit threshold at which investments take place. There is a threshold for investment in the shim layer information cloud that is lower than for investment in the native information cloud, with the cost of transformation of the network to native information cloud also being higher. If the investment for both types of transformation is set to very high, then we effectively have a scenario that is ‘IP only’ in which no network transformation takes place. If only the native information cloud traffic threshold is high, we refer to this scenario as ‘shim invest’ as the threshold for investment in the native information cloud results in little or no network capability for native traffic. If both thresholds are low, then this scenario is labelled as ‘shim native invest’ as investments in both types of network will take place within the timescale of the model.

In addition to the dimension of willingness to invest, there is a strong determining factor according to the relative cost of implementation, compared with the profitability of Information Cloud traffic. It is fairly easy to predict that if the network changes are inexpensive to make, that this will have a highly beneficial effect on uptake, but it is interesting to consider the sensitivity of this relative financial benefit to the uptake of the Information Cloud over 20 years, so this is the other main variable that has been varied in carrying out an exploration of scenarios.

7.8.1 Relative Cost of Network Transformation

First we consider the effect of the relative cost of transformation on the Total Profit achieved by a network operator, over 20 years, based on three different network transformation costs: “cheap”, “expensive” and standard. All other variables remain constant. The investment levels in terms of network transformation and application sectors are all high, or ‘proactive’. The relative costs of network transformation are given in the following graphs:

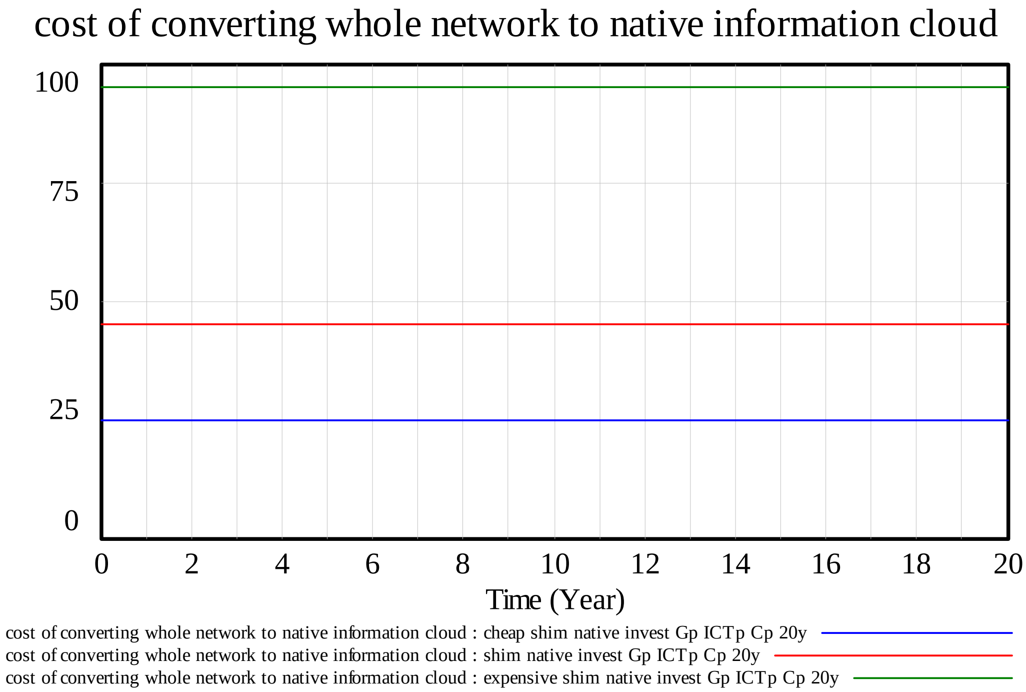

Figure 7.5: Relative costs in network transformation to native information cloud

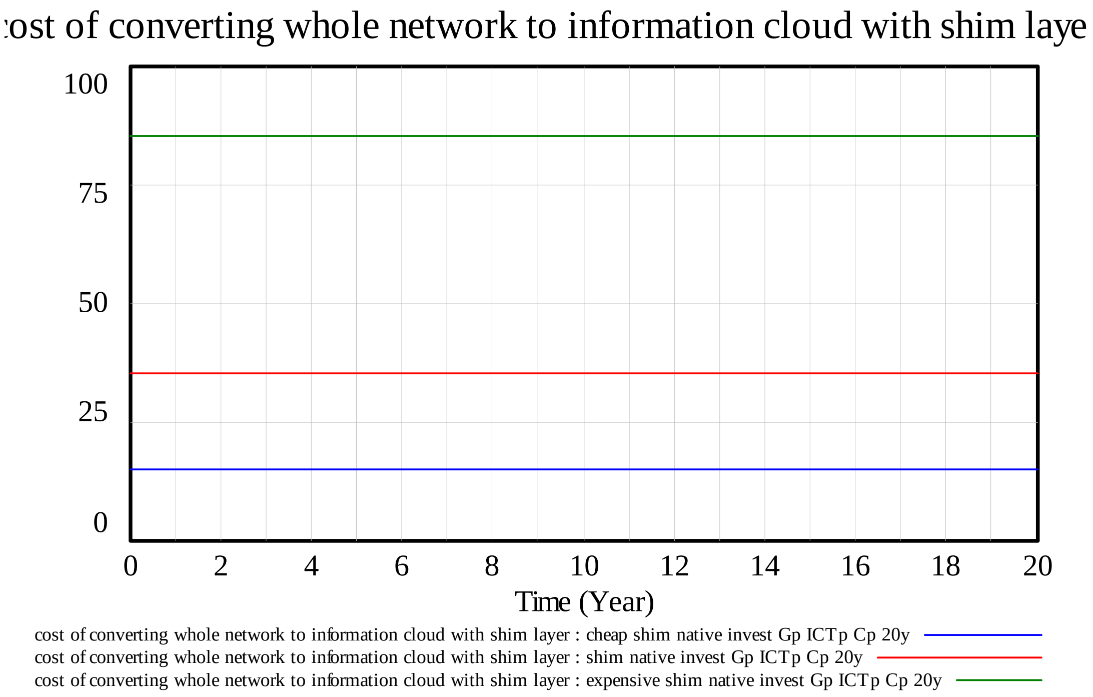

Figure 7.6: Relative cost in network transformation to ‘shim layer’ Information Cloud

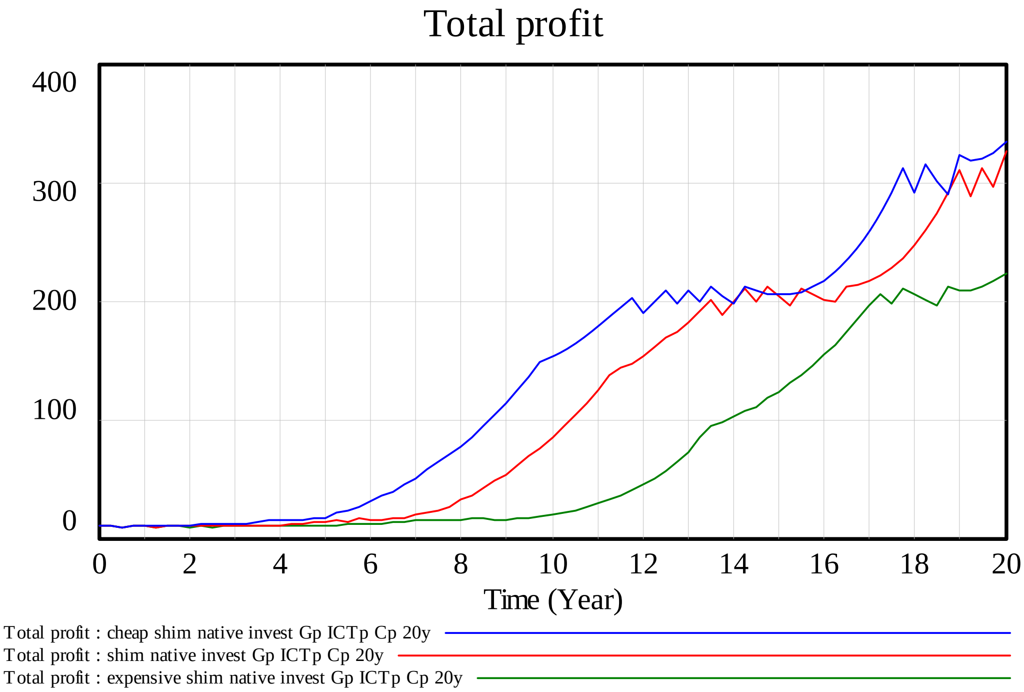

Figure 7.7: Relative profit levels according to the cost of network transformation

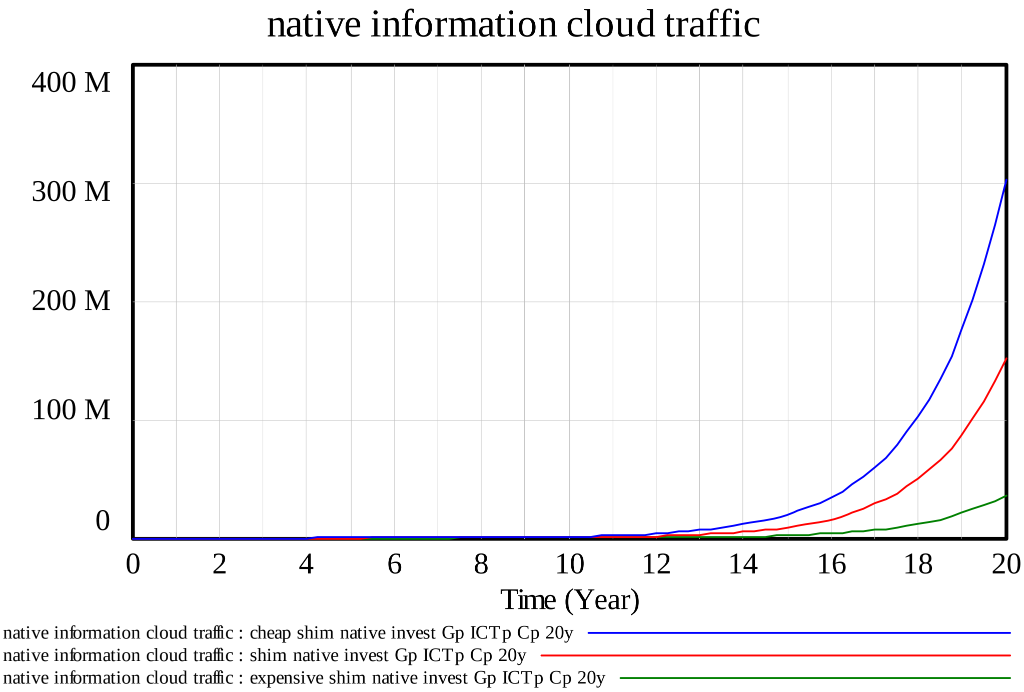

Figure 7.8: Dependence of ‘native’ Information Cloud traffic on the cost of network transformation

We see in Figure 7.7 that when the cost of network transformation is reduced, profit levels rise more rapidly, as would be expected. Note, however, that the profit levels do not simply follow the traffic levels, which show an exponential response. Regulator or market drivers are modelled that reduce the unit price (revenue) of traffic at levels above 1M units and then again above 100M units. In addition, we see that there is a 'plateau' in profits around the level of 200 and then around the level of 300. This is a consequence of also modelling the ability of the network operator to invest further in the network in order to reduce the unit cost of handling traffic, as shown in the following graphs:

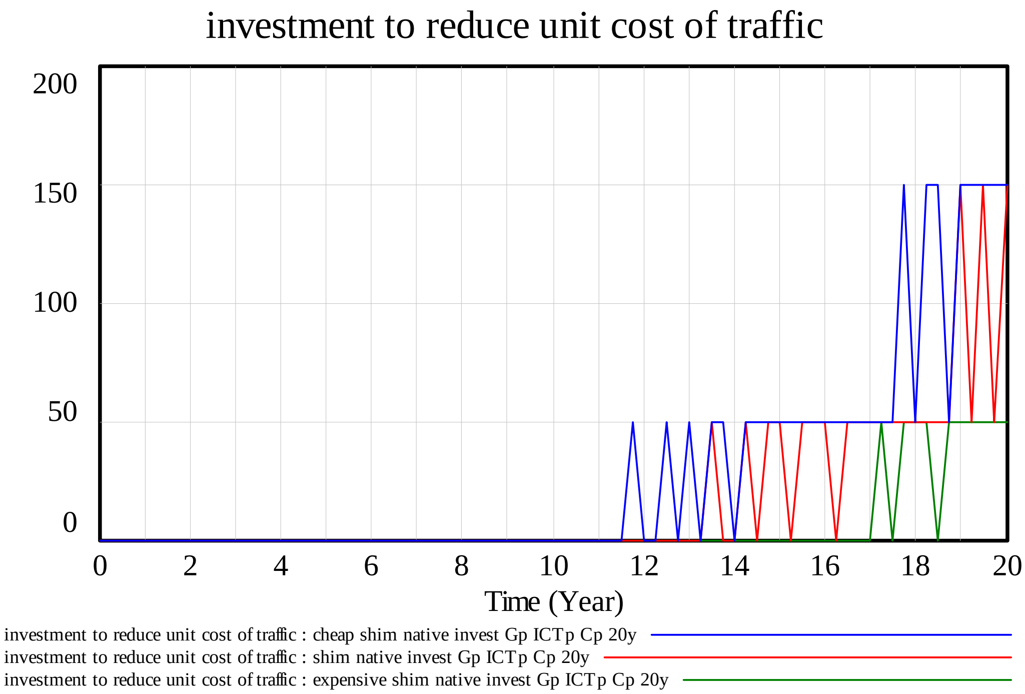

Figure 7.9: Investment by the network operator to reduce the unit cost of forwarding traffic

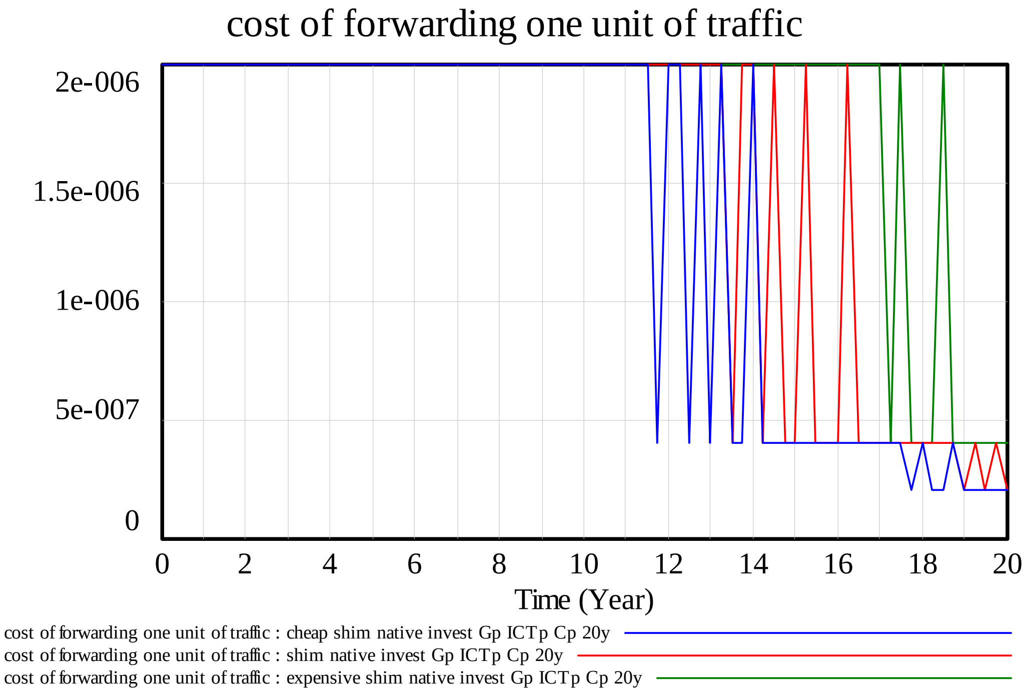

Figure 7.10: How the cost of forwarding a unit of traffic varies depending on the cost of network transformation, allowing for network investment

The conclusion from this is that evidently lowering the costs of network transformation will have a large impact on the rate of generation of new traffic types. However, under a regulatory regime and market forces that would not stand exponential revenue growth, and bearing in mind the need to invest in lowering the unit cost of handling traffic, in order to meet those much higher traffic volumes, the Total profit levels are less sensitive to the cost of network transformation than the traffic volumes.

7.8.2 Effect of Investment Decisions

We next consider the effects of investment, both by the network operator and by application sectors, on the migration to the Information Cloud.

In the following graphs, the label ‘IP only’ refers to Total Profit thresholds in order to invest in shim layer and native information cloud as being 100 and 150 respectively. In other words, these are very high and may never be reached. The label ‘shim invest’ implies Total Profit thresholds in order to invest in shim layer and native information cloud as being 10 and 150, meaning that shim layer investment will take place but native information cloud investment will be very slow or absent. The label ‘shim native invest’ implies Total Profit thresholds in order to invest in shim layer and native information cloud as being 10 and 15, which means the network operator is proactive in investing early in both types of network. The labels G, ICT and C, refer to Government, Collaborative ICT and Content-Centric investment respectively. The subscript ‘s’ refers to ‘sceptical’ with a low level of constant investment. The subscript ‘p’ refers to ‘proactive’ which corresponds to a high level of constant investment.

Further detail on those investment levels is given below:

> Gs = 100 + 0.05(information cloud shim layer traffic), for information cloud shim layer traffic <1e+007, above which investment is capped at 55,000 per annum
>
> Gp = 6000 + 0.05(information cloud shim layer traffic), for information cloud shim layer traffic <1e+007, above which investment is capped at 55,000 per annum

ICTs = 100 + 0.03(information cloud shim layer traffic) + 0.05(native information cloud traffic), for native information cloud traffic <2e+007, above which investment is capped at 100000 per annum

ICTp = 6000 + 0.03(information cloud shim layer traffic) + 0.05(native information cloud traffic), for native information cloud traffic <2e+007, above which investment is capped at 100000 per annum

Cs = 100 + 0.01(information cloud shim layer traffic) + 0.015(native information cloud traffic), for native information cloud traffic < 1e+006, above which investment is capped at 150000 per annum

Cp = 1000 + 0.01(information cloud shim layer traffic) + 0.015(native information cloud traffic), for native information cloud traffic < 1e+006, above which investment is capped at 150000 per annum

Note that there are assumptions built into the way that investment levels vary with information cloud traffic levels. For example, there is no dependence of Government investment on ‘native information cloud traffic’ as the required functionalities are not thought to depend primarily on native PSIRP solutions. Another major assumption is that a proactive content-centric sector will have a relatively smaller constant level of investment (i.e. traffic independent) as it is assumed that the success of content-centric applications, and therefore the feed-forward for further investment, will be dominated by traffic volumes. By contrast, it is assumed that the Government and Collaborative ICT sectors would be willing to make more significant traffic-independent investments to achieve their information cloud goals, because their primary goals can be met from single-tenanted solutions, or at least without requiring mass market use of the Information Cloud.

Eleven scenarios were modelled: “IP only”, “shim invest Gs ICTs Cs”, “shim invest Gp ICTs Cs”, “shim invest Gs ICTp Cs”, “shim invest Gs ICTs Cp”, “shim invest Gp ICTp Cp”, “shim native invest Gs ICTs Cs”, “shim native invest Gp ICTs Cs”, “shim native invest Gs ICTp Cs”, “shim native invest Gs ICTs Cp” and finally “shim native invest Gp ICTp Cp” in which all investment is proactive.

These scenarios enable us to consider the importance of concurrent or near-concurrent investment in both information cloud solutions, and to consider if one application sector investment is more important than another in achieving the network transformation.

7.9 Model Results Over 20 Years

In Figure 7.11, we see the total traffic growth over 20 years for the eleven scenarios outlined. We see that the greatest traffic volume is for the scenario with proactive investment in all three sectors and by the network operator. The next two highest traffic volumes are for scenarios with proactive network investment. However, the scenario with proactive Collaborative ICT only does better than the scenario with proactive Government investment only. We note that proactive network investment with proactive Content-Centric investment alone is slightly worse than shim layer investment only, when combined with proactive investment in all three application sectors.

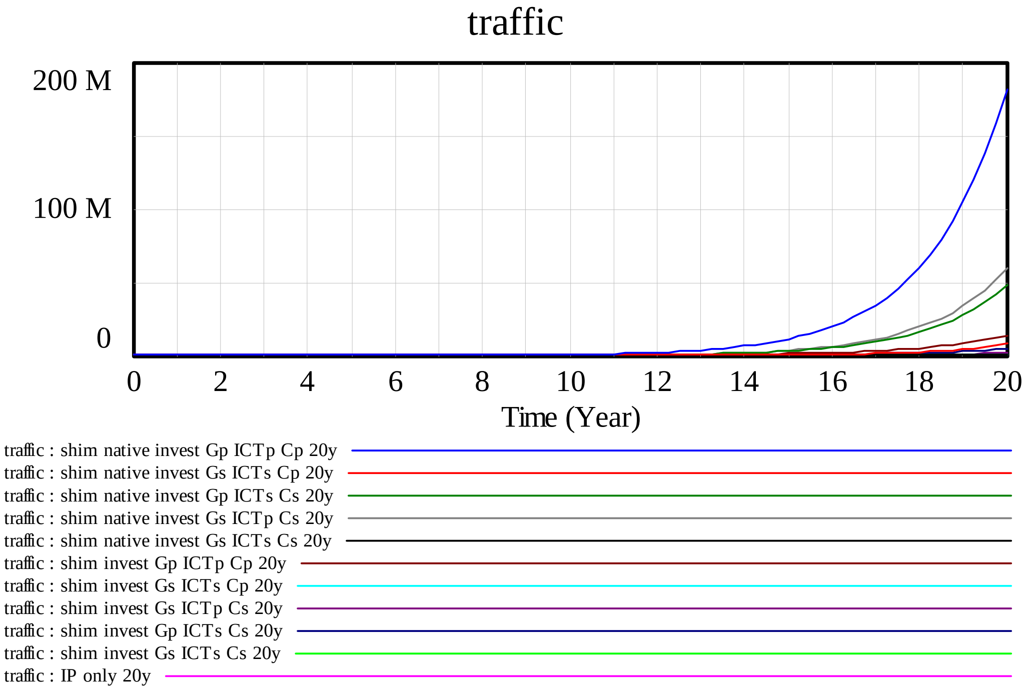

Figure 7.11: Growth of network traffic for sub-scenarios over 20 years

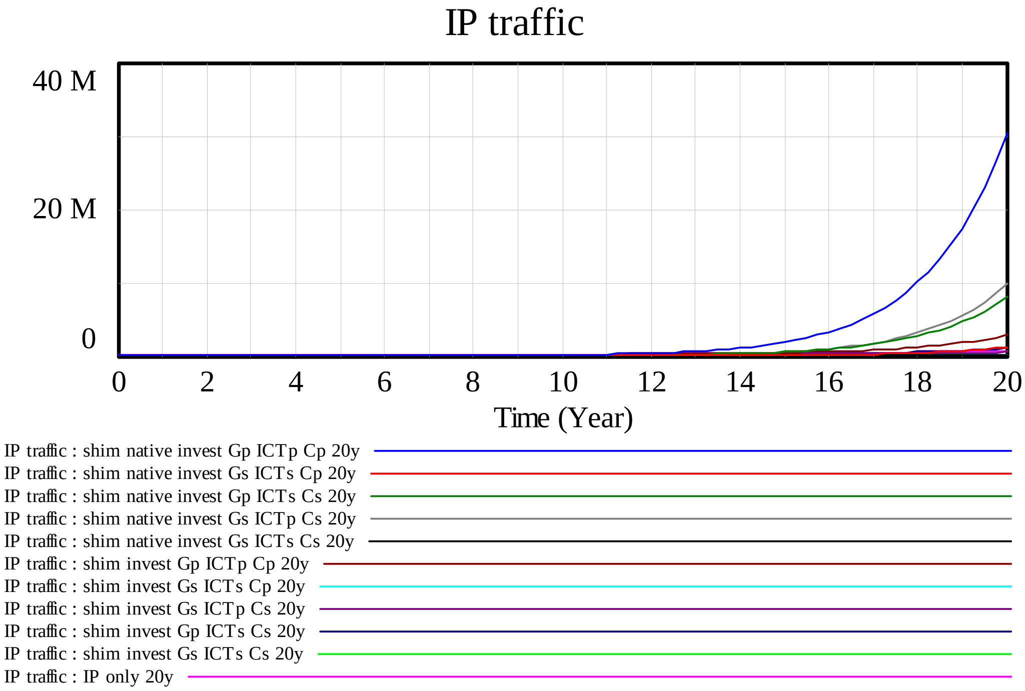

Figure 7.12: Growth in IP traffic for sub-scenarios over 20 years

In Figure 7.12, we see that what is termed ‘IP traffic’ grows along with the information cloud traffic types. This is modelled to reflect the fact that growth in network traffic will continue independently of the Information Cloud. However, it seems likely that once the network is fully converted to the native Information Cloud, that most or all of what we have called ‘IP traffic’ will also be converted to the Information Cloud.

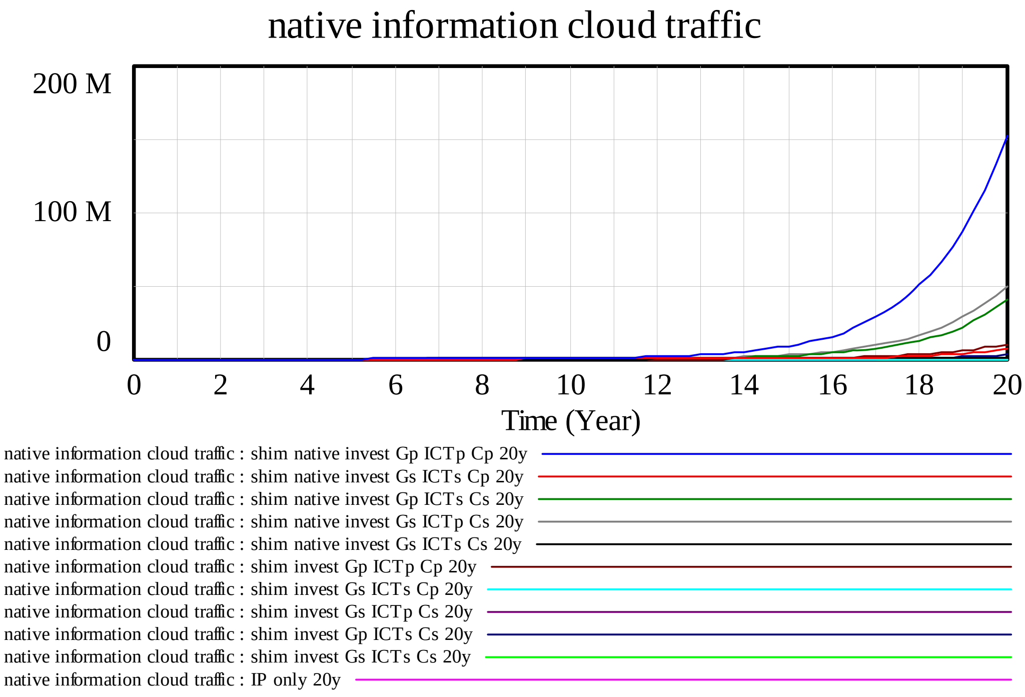

Figure 7.13: Growth in ‘native’ Information Cloud traffic over 20 years

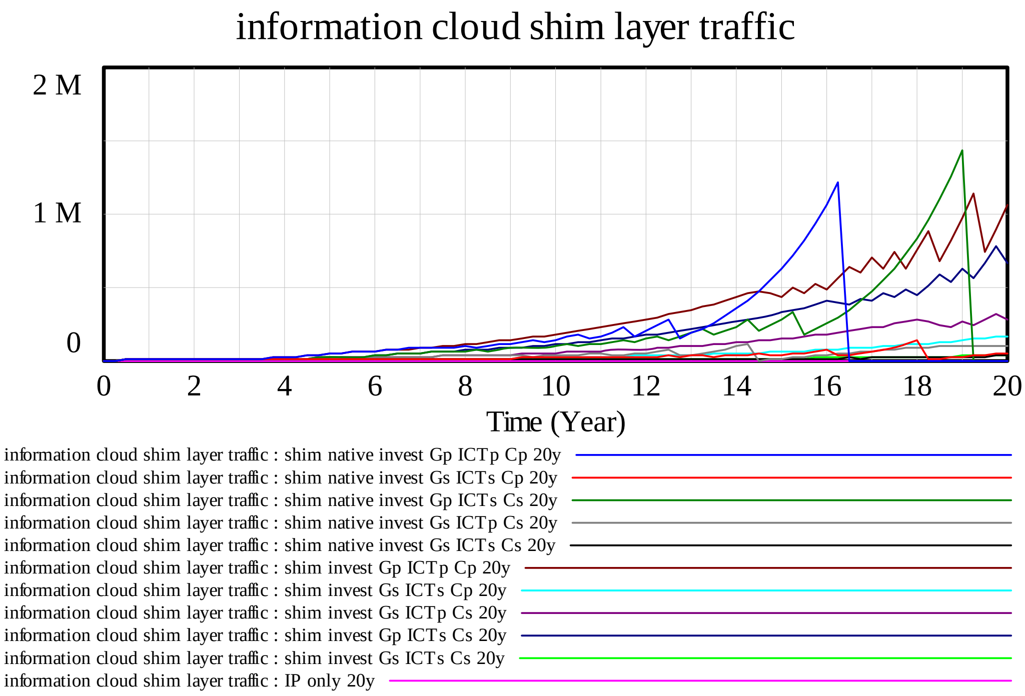

Figure 7.14: Quantity of ‘shim layer’ Information Cloud traffic over 20 years

We see in Figure 7.14 that shim layer traffic tends to decline for the scenarios where there is proactive native information cloud investment. That is because the model assumes that shim layer traffic will be converted into native traffic, once there is sufficient native Information Cloud capacity.

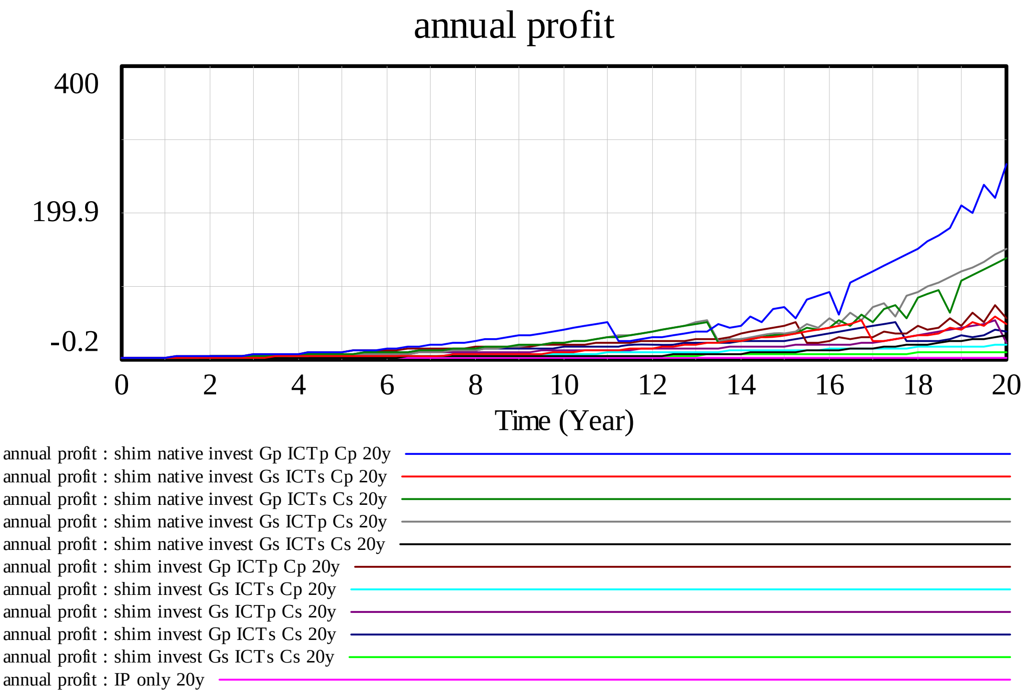

Figure 7.15: Annual profit for the network operator

In Figure 7.15, we see that the markedly higher profit levels are for the scenarios with proactive network investment and where there is proactive investment from Government or Collaborative ICT or all three sectors. Profit levels with only application investment from the Content-Centric sector perform significantly worse, especially when the network investment is also more cautious, with a high threshold for investment in native information cloud, as we see in the turquoise line (shim invest Gs ICTs Cp 20y.)

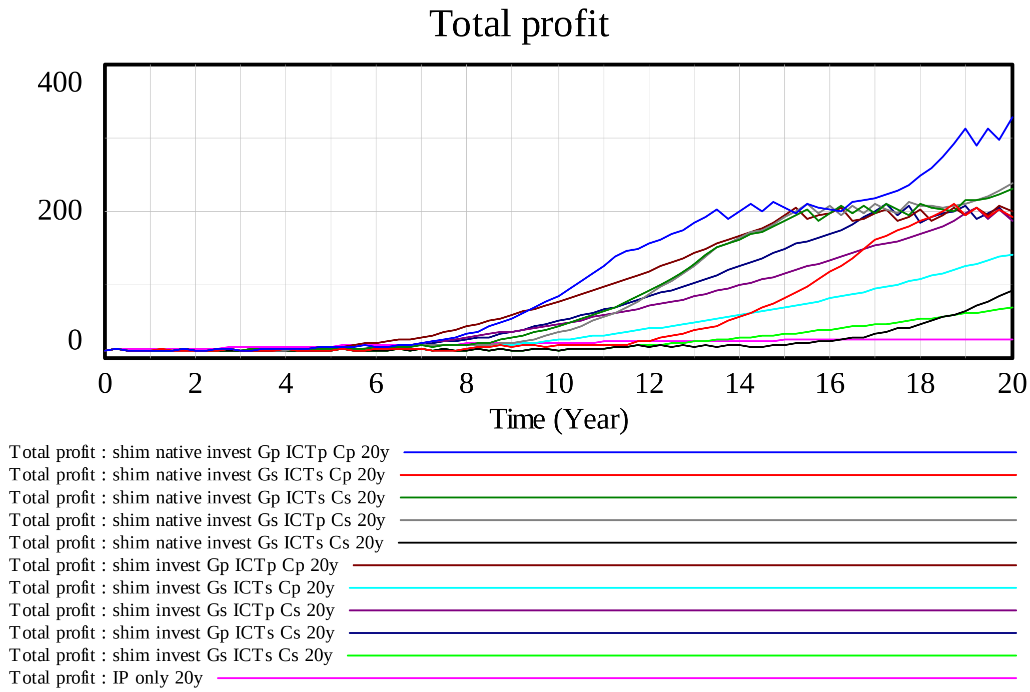

Figure 7.16: Cumulative profit for the network operator

In Figure 7.16, we see that Total profit levels do not follow the same curves as for traffic volumes. This is a consequence of the way that profit has been modelled, with the expectation that revenues will not grow anywhere near proportionally with traffic volume. Additionally, traffic volume costs are modelled as rising linearly, which means that for very high traffic volumes, losses could be made unless this unit cost is reduced. Therefore, we model investment rates for the reduction in unit cost of network traffic, which slows the growth in Total profit. If this modelling were accurate, it would mean that over a 16 year timescale the top three performing scenarios have around the same Total profit levels. However, over a longer timescale, the scenario with all three proactive sectors diverges from the closest performing scenarios, which have proactive investment in either Government or Collaborative ICT only. However, it is interesting to note that proactive investment in either of these would create compelling circumstances for proactive network investment, according to this model.

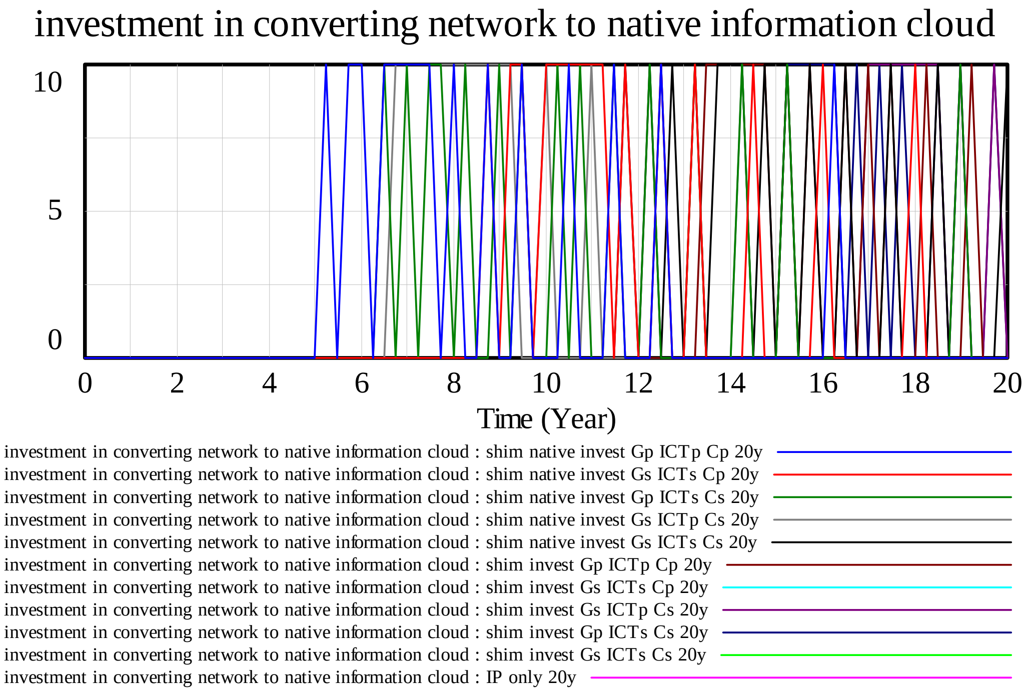

Figure 7.17: Network investment in transformation to 'native' Information Cloud

In Figures 7.17 and 7.18, we see that the scenario with proactive network and sector investments (blue line) makes the earliest investments in network transformation.

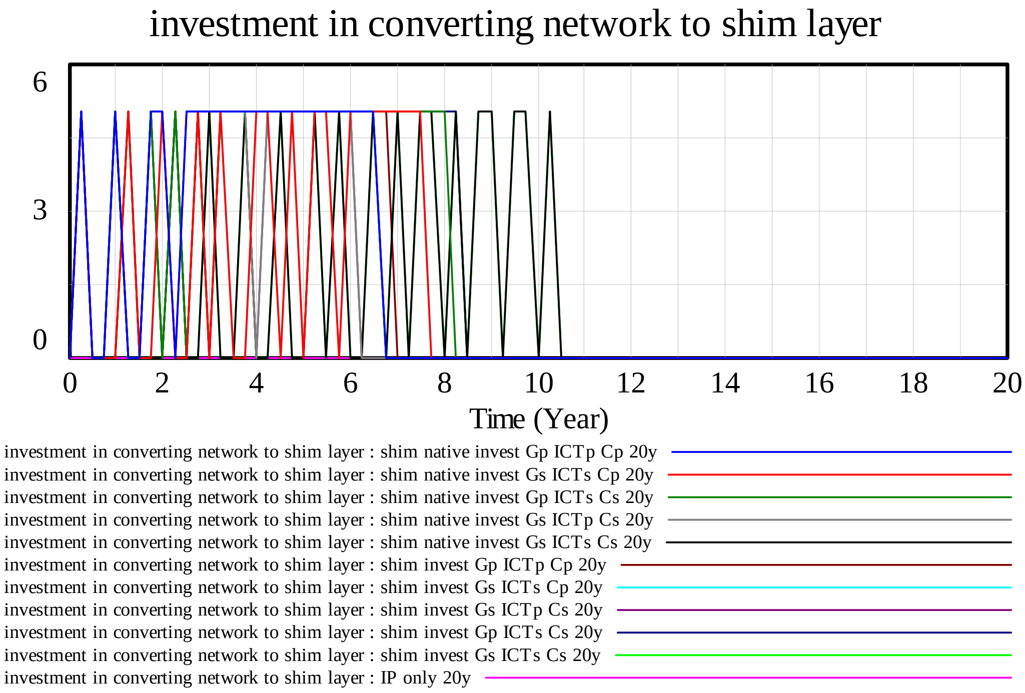

Figure 7.18: Network investment in transformation to 'shim layer' Information Cloud

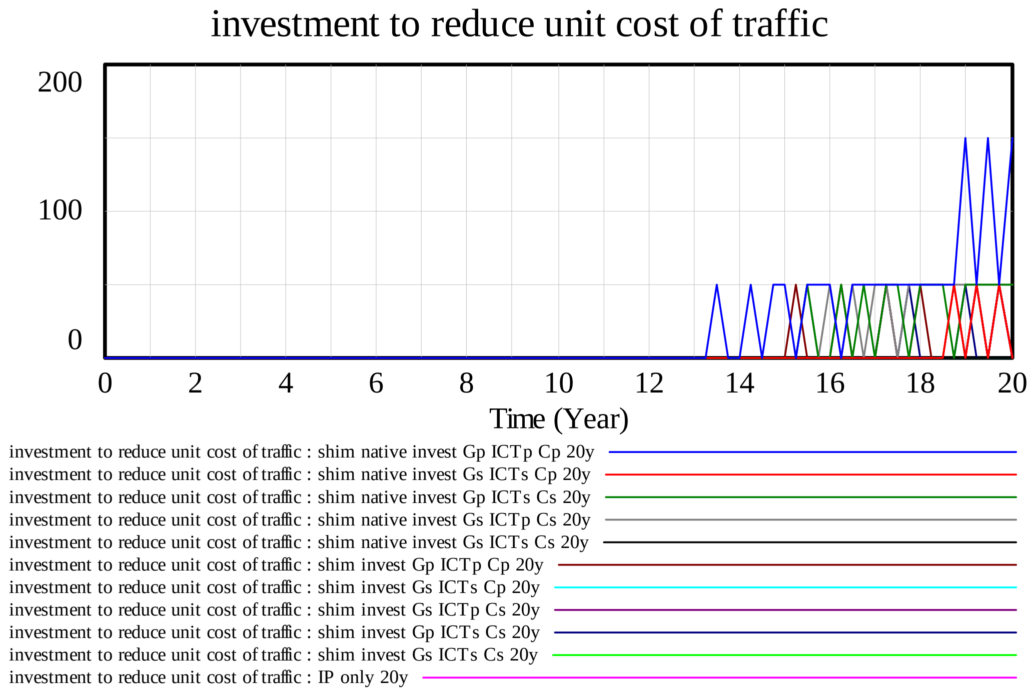

Figure 7.19: Network investments to reduce unit traffic cost

In Figure 7.19, we see that the scenario with the proactive network and sector investments (blue line) is the only one which, within the 20 year timescale, has begun to make increased investment to reduce the unit cost of traffic, in response to the highest traffic volumes.

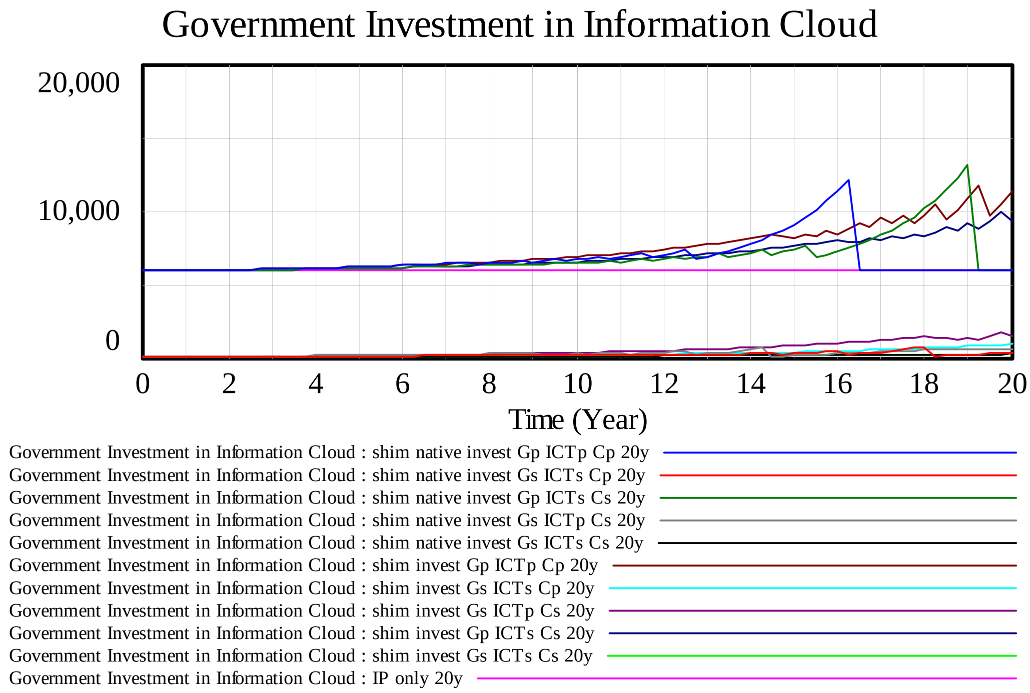

Figure 7.20: Government sector investment in Information Cloud applications

In Figure 7.20, we see that Government investment in the Information Cloud can fall back to starting levels. This is because the model links this investment to shim layer traffic volumes only, which can decline to almost zero, as native information cloud traffic takes over. It is reasonable to expect that, in investment terms, a government might be willing to be an early adopter of a new technology, but would then have finite revenue drivers, which would constrain further investment, once all its operations and services had been transformed.

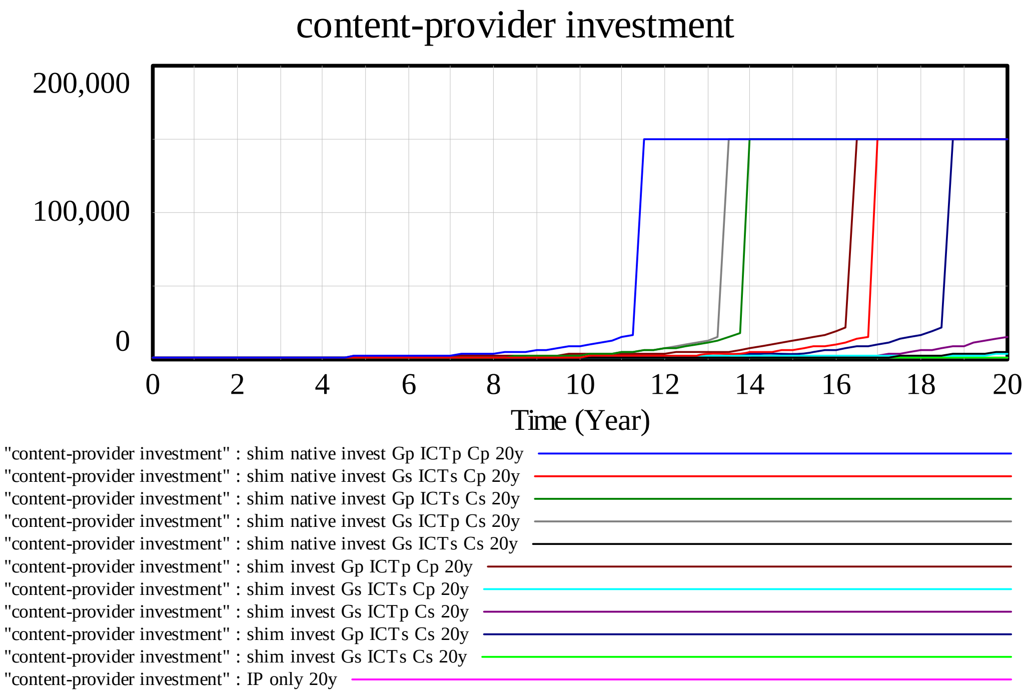

Figure 7.21: Investments from the content-centric sector in Information Cloud applications

In Figure 7.21, we see that Content-provider investment fairly rapidly achieves its maximum level, at which it stays, within the model. The provision of content is not expected to require exponential investment, as its audience grows linearly.

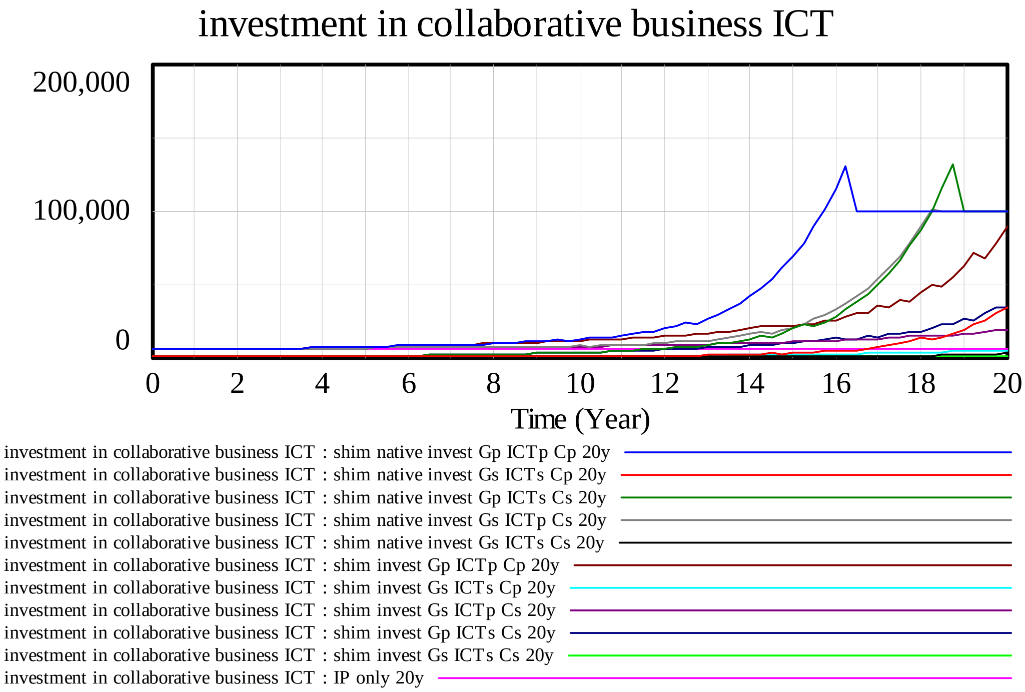

Figure 7.22: Investments from the Collaborative Business ICT sector in Information Cloud applications

In Figure 7.22, we see that investment in Collaborative ICT can fall back from a peak. The model imposes a cap on the level of Collaborative ICT investment, beyond which it does not increase with increased traffic volumes. Again this is to recognize that investment will reach a plateau, once all businesses are transformed.

7.10 Model Results Over 10 Years

The following figures present the results from the same scenarios, over the first ten years, for the sake of clarity.

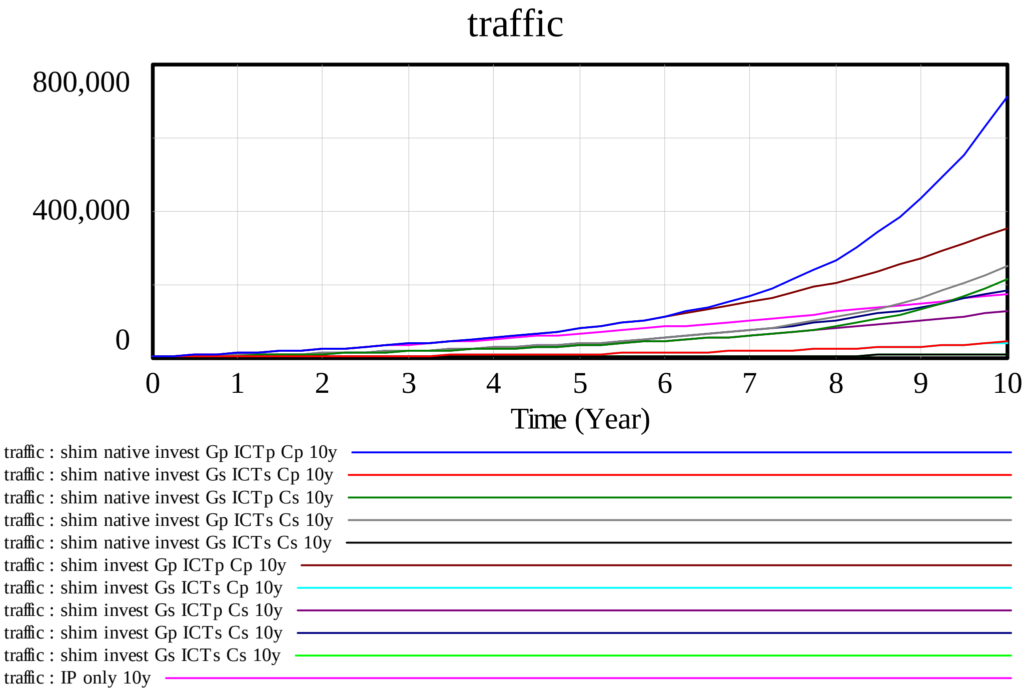

Figure 7:23 Total traffic volumes, for the sub-scenarios, over 10 years

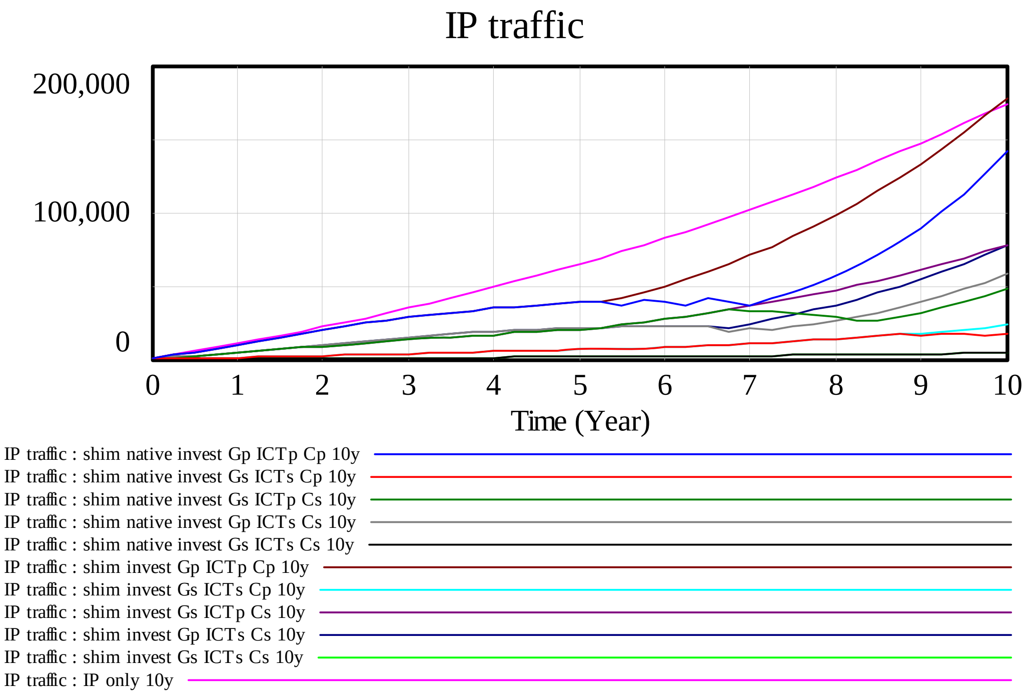

Figure 7.24: IP traffic volumes over 10 years

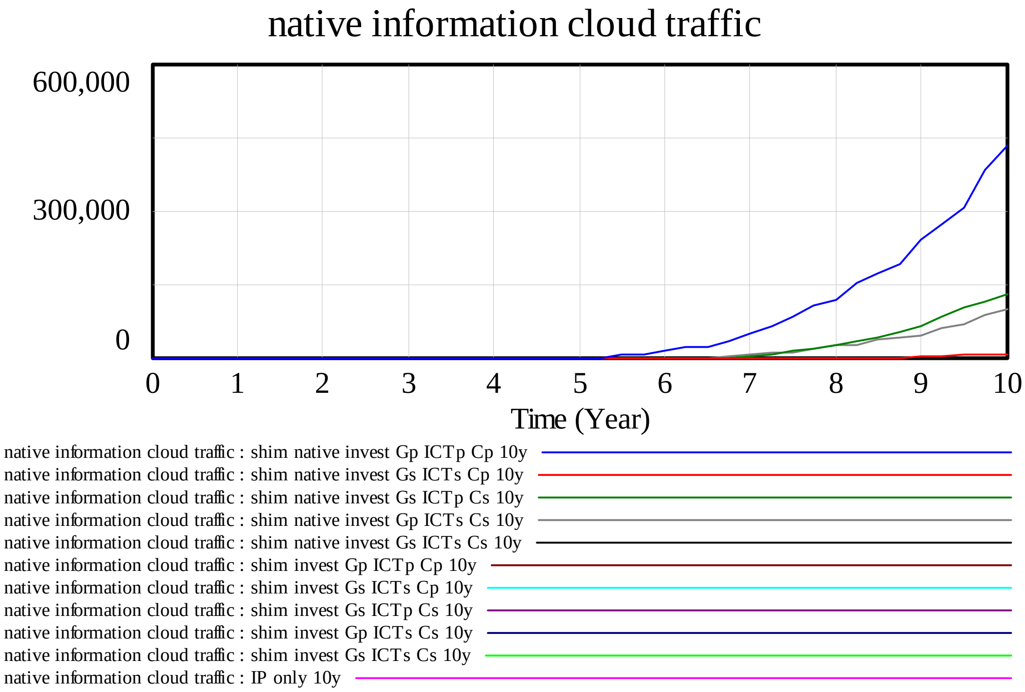

Figure 7.25: 'Native' Information Cloud traffic volumes over 10 years

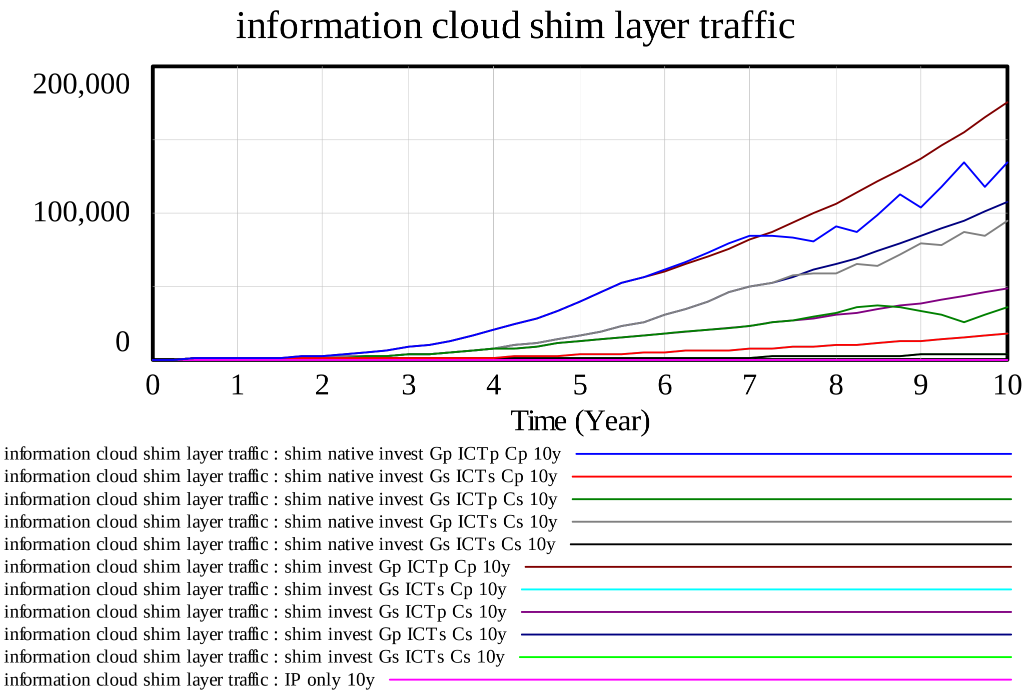

Figure 7.26: 'Shim layer' Information Cloud traffic volumes over 10 years

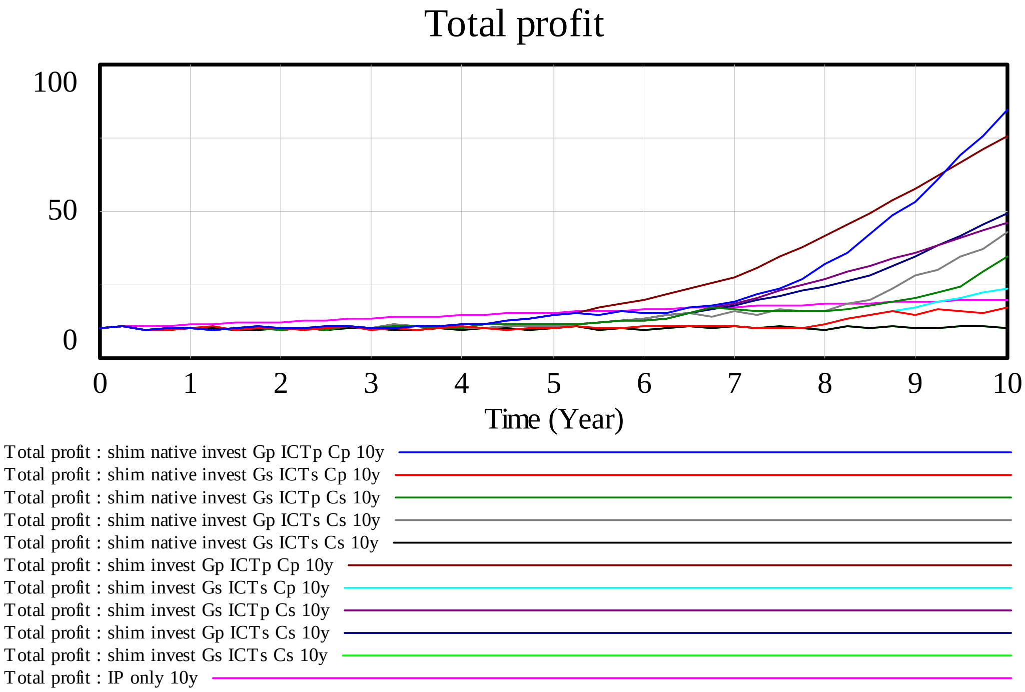

Figure 7.27: Cumulative profit levels for the network operator, over 10 years

We see in Figure 7.27 that the scenarios in which all application sector investment is proactive, give the highest Total profit levels. However, we see that for predominantly shim investment (brown line), as compared with shim and early native investment (blue line) the profits are higher until after 9 years into the scenario. Other shim native scenarios are even slower to reach the best Total profit levels. This is because Total profit has subtracted from it the cost of network investments. With greater investment in native Information Cloud transformation, the higher profits are barely or not at all seen over a ten year timescale. The higher returns come over longer timescales, depending on the levels of sector investment.

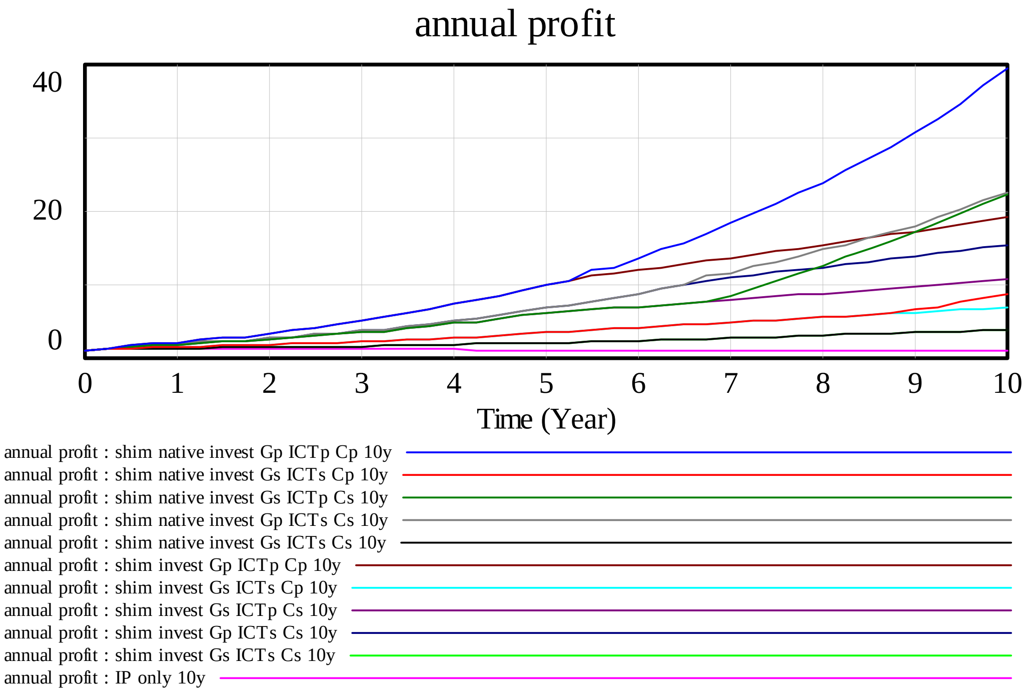

Figure 7.28: Annual profit levels for the network operator, over 10 years

In Figure 7.28, we see the annual profit levels for the shim native investment with proactive application sectors are the highest, as these values are calculated before network investments, and refer to revenue, less traffic costs.

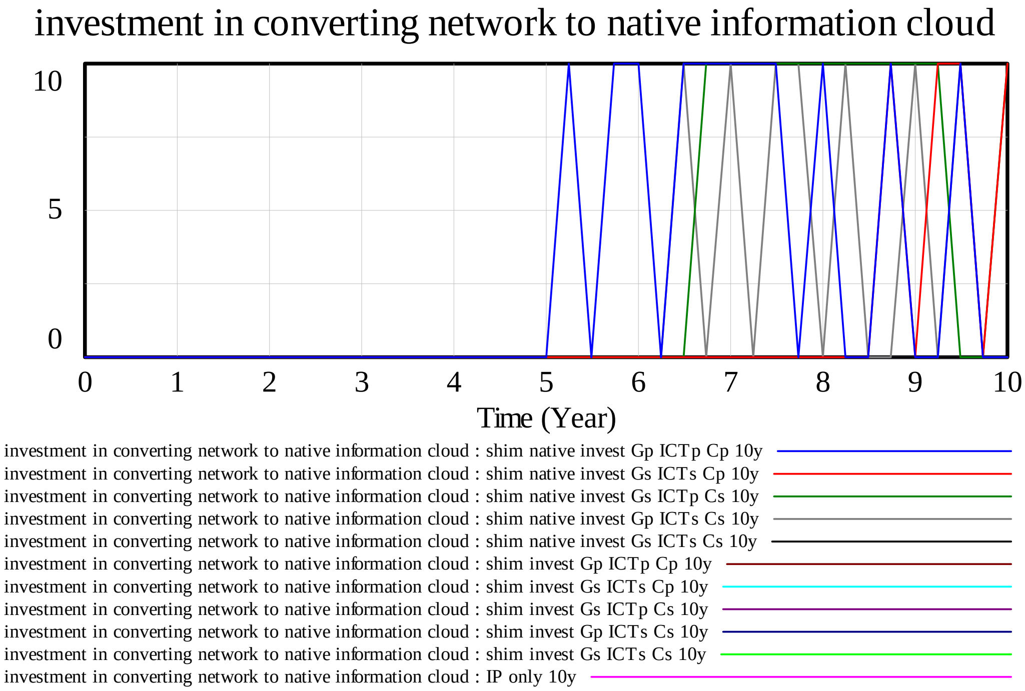

Figure 7.29: Network investments in transformation to 'native' Information Cloud

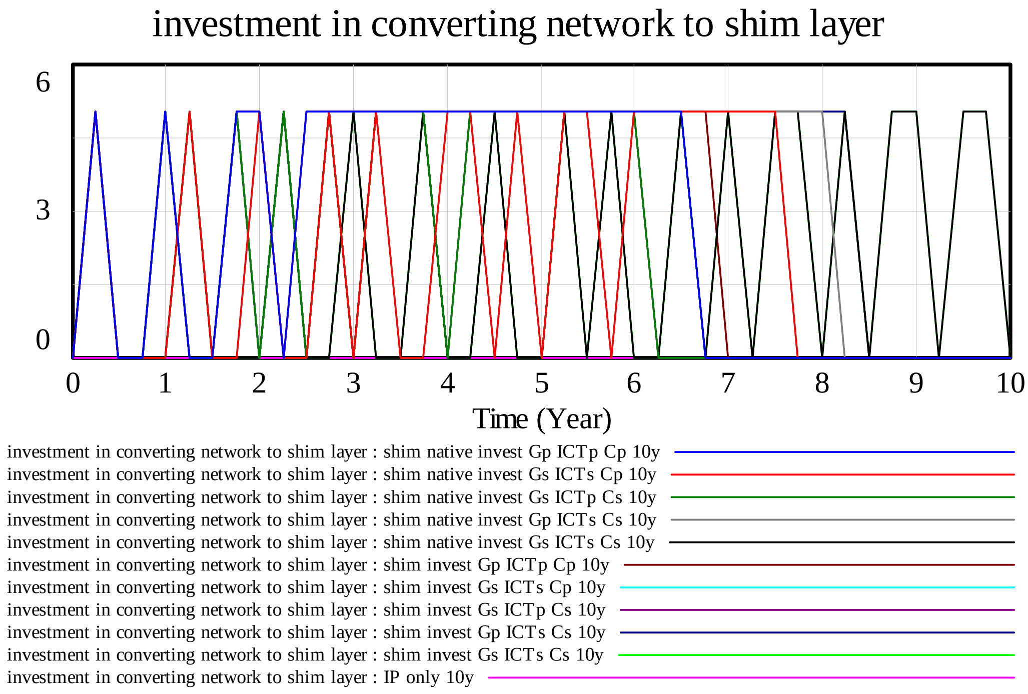

Figure 7.30: Network investment in transformation to 'shim layer' Information Cloud

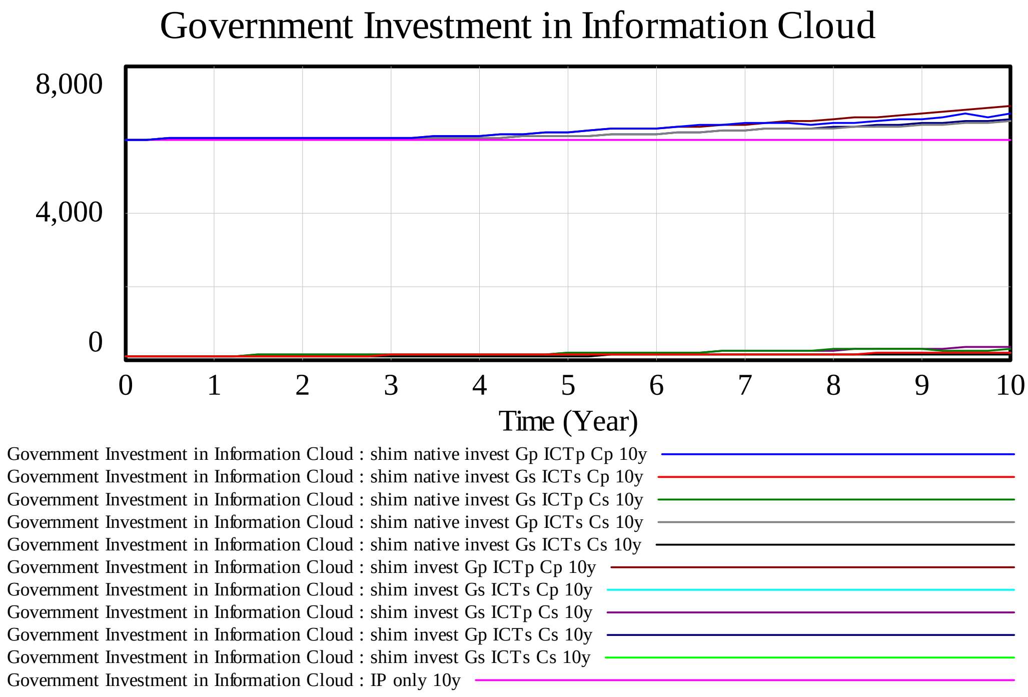

Figure 7.31: Government sector investment in Information Cloud applications, over 10 years

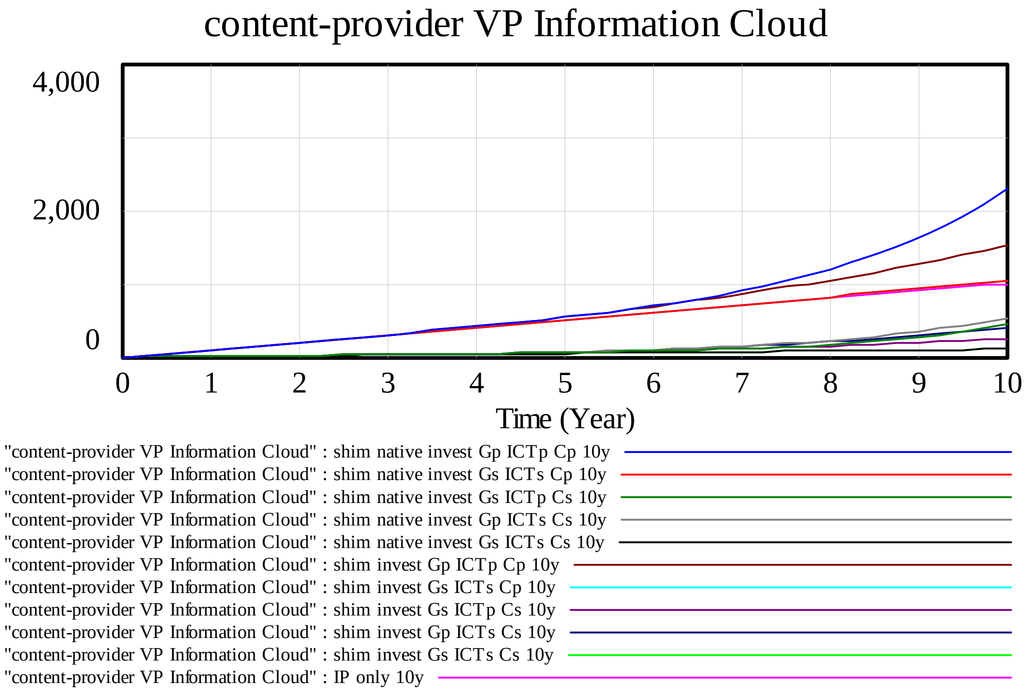

Figure 7.32: Investment from the Content-Centric sector in Virtual Private Information Cloud solutions

In Figure 7.32 we see that investment from the Content-Centric sector causes growth in Virtual Private Information Cloud for content-providers. This in turn stimulates demand for Information Cloud traffic.

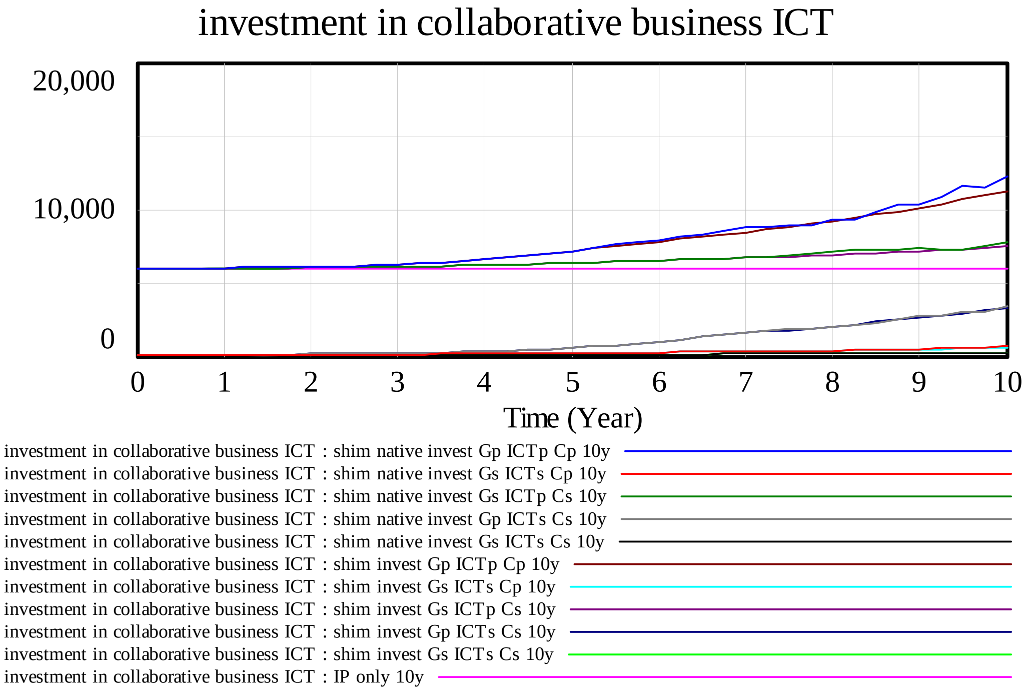

Figure 7.33: Investment by the Collaborative Business ICT sector in Information Cloud applications

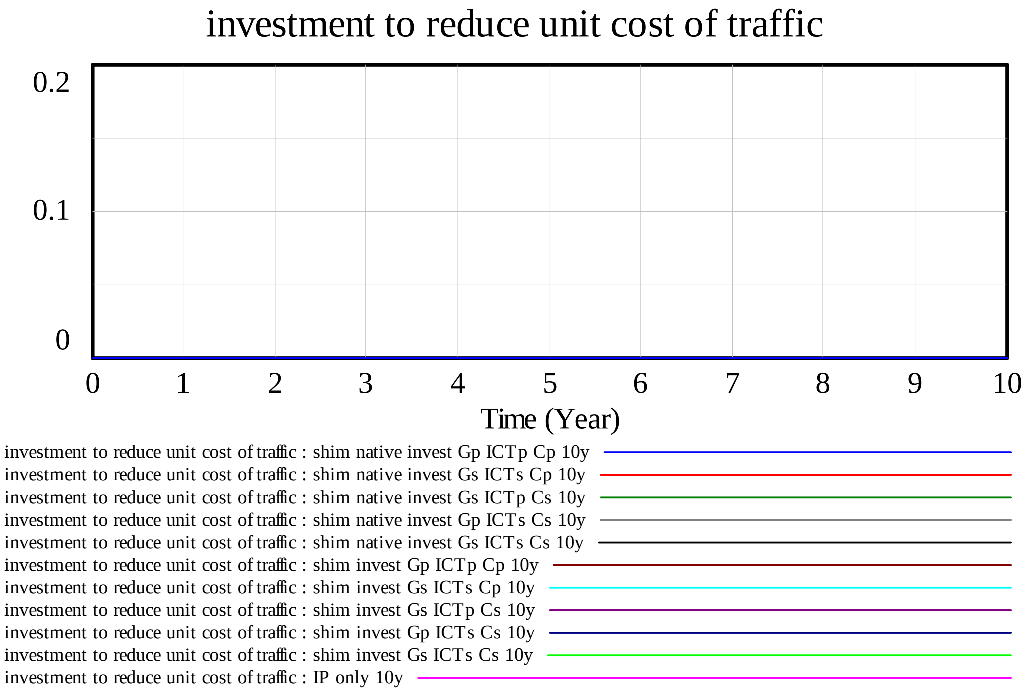

Figure 7.34: Network investment to reduce traffic in the first 10 years

Traffic levels, for these scenarios, do not reach sufficiently high levels, within the first ten years of simulation, to require the network operator to make additional investment to reduce the unit cost of forwarding traffic.

7.11 Model Results Over 30 Years

In the following figures, we see the same scenarios, but left to run for thirty years. The results are very similar, but we see the traffic volumes soar.

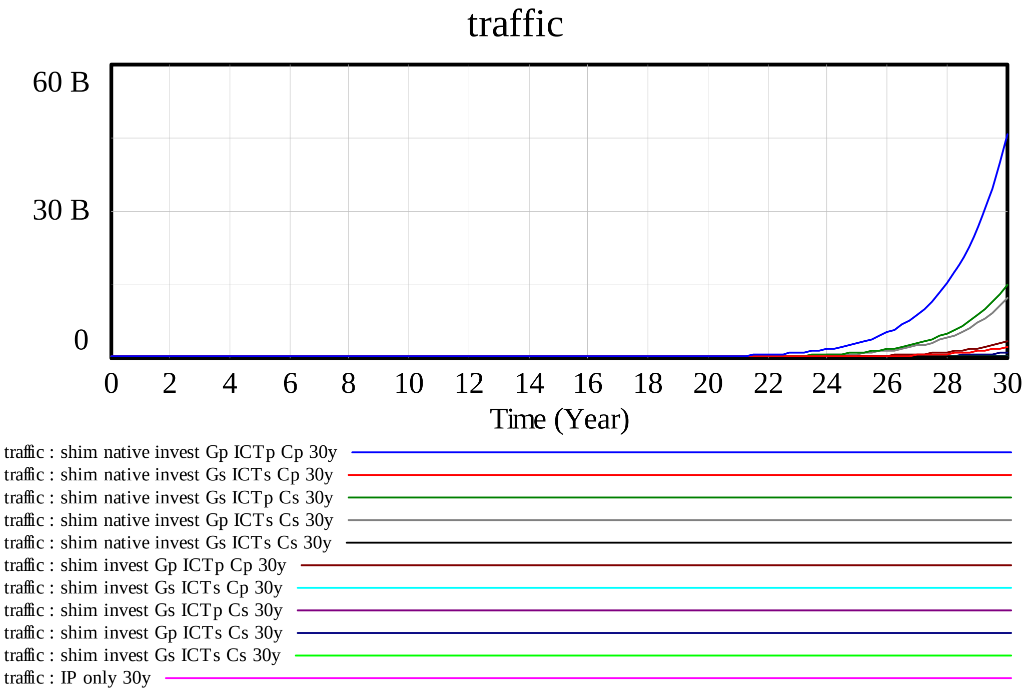

Figure 7.35: Traffic volumes over 30 years

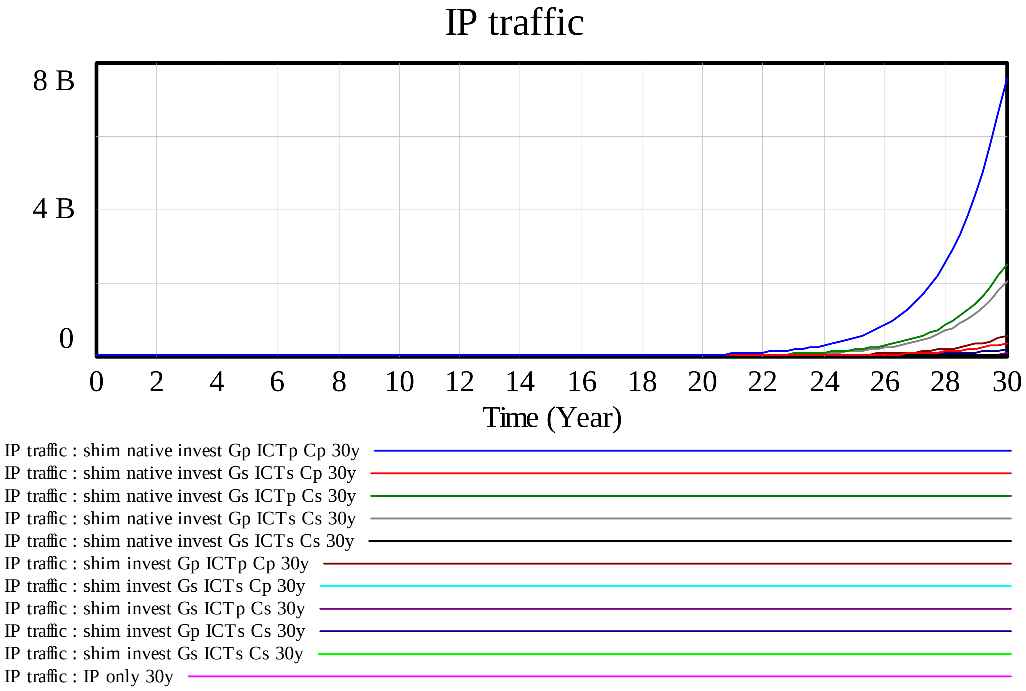

Figure 7.36: IP traffic volumes over 30 years

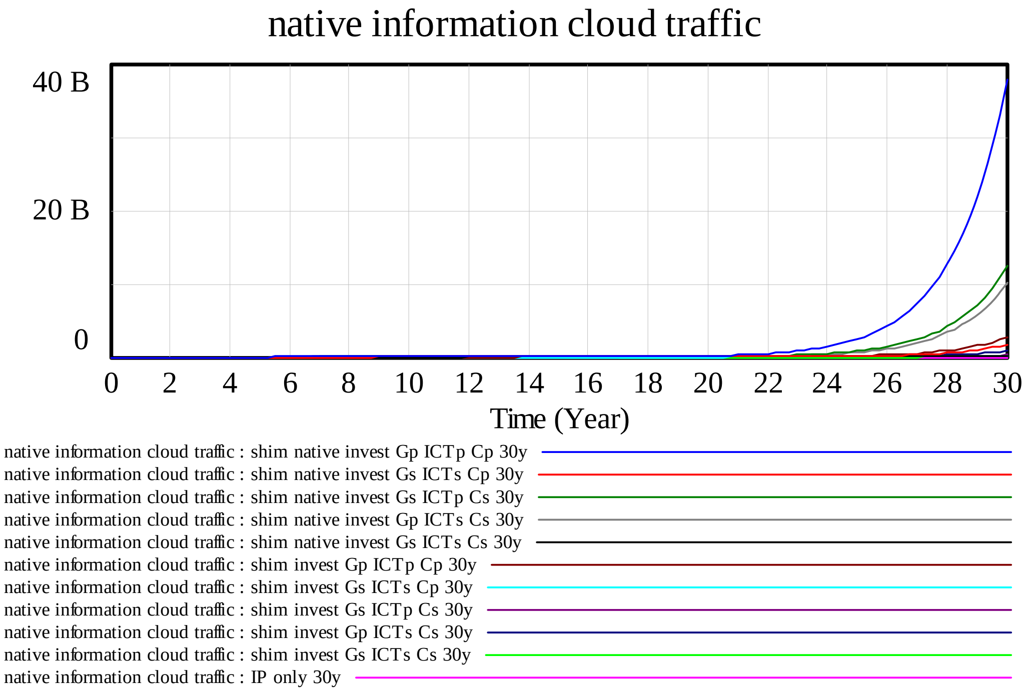

Figure 7.37: ‘Native’ Information Cloud traffic volumes over 30 years

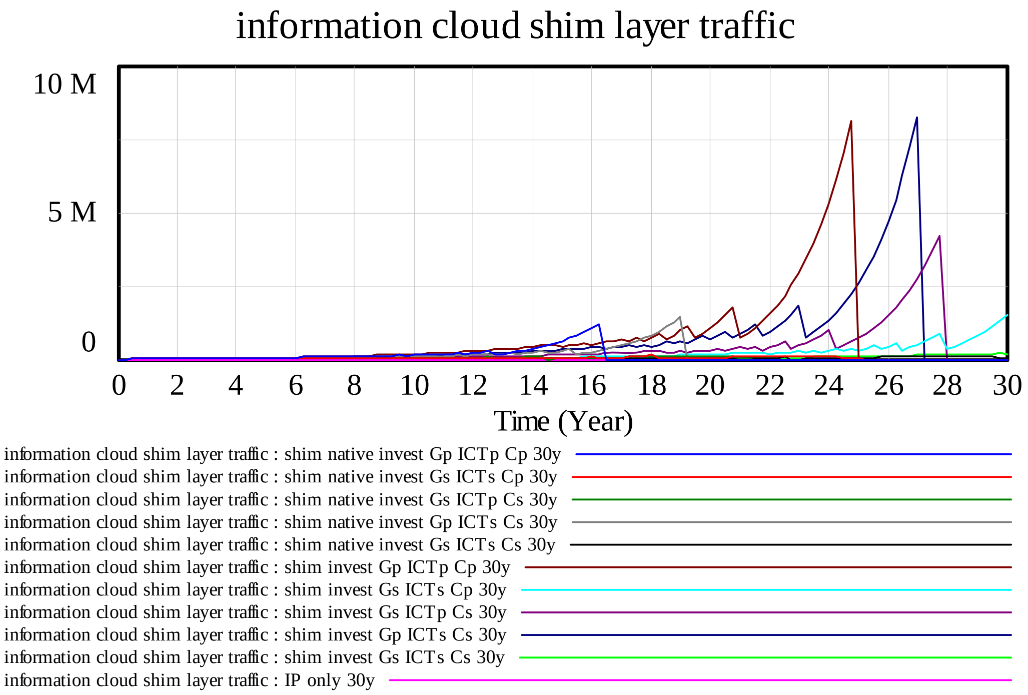

Figure 7.38: ‘Shim layer’ Information Cloud traffic volumes over 30 years

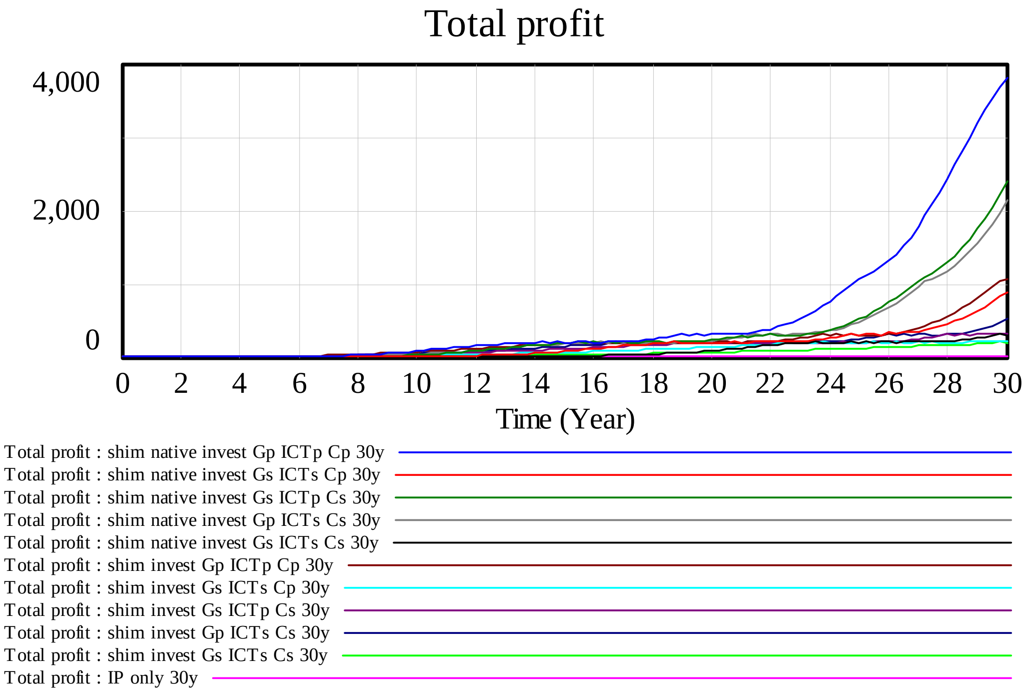

Figure 7.39: Cumulative profit for the network operator over 30 years

In our model, 'Total profit' levels increase dramatically between 20 an 30 years of the simulation. However, if we look at Figure 7.40 we see that 'Annual profit' levels are starting to fall for the most successful proactive investment scenario. This is because traffic volumes are rising so fast that the cost of handling that traffic in the network is beginning to outweigh the revenues from the traffic. Therefore, if the model were accurate, the network operator would be required to make another network investment towards the end of the 30 year period, in order to reduce traffic costs, to safeguard their profits, going forward. However, this is sufficiently far into the future that a new technology and network paradigm may have appeared by this time.

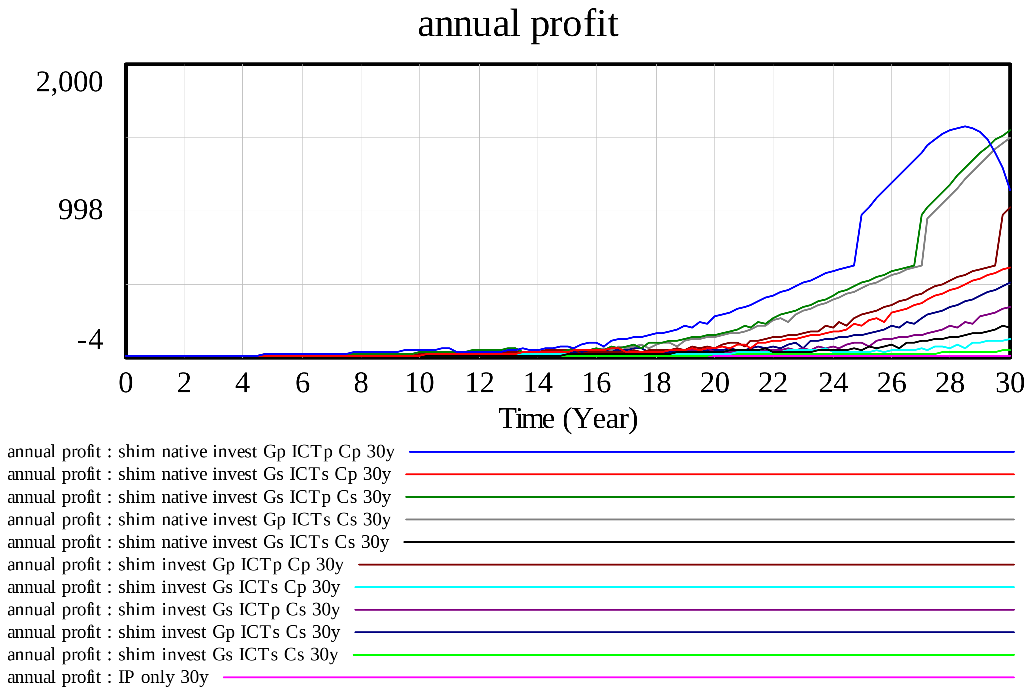

Figure 7.40: Annual profit levels for the network operator, before network investment costs

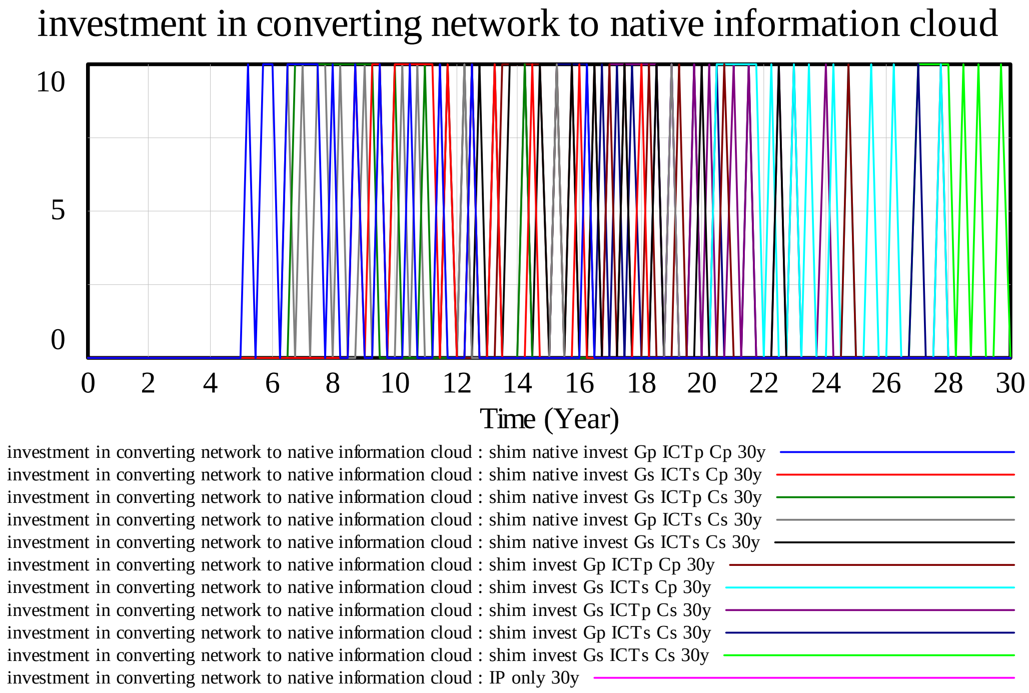

Figure 7.41: Network investments in transformation to a 'native' Information Cloud architecture

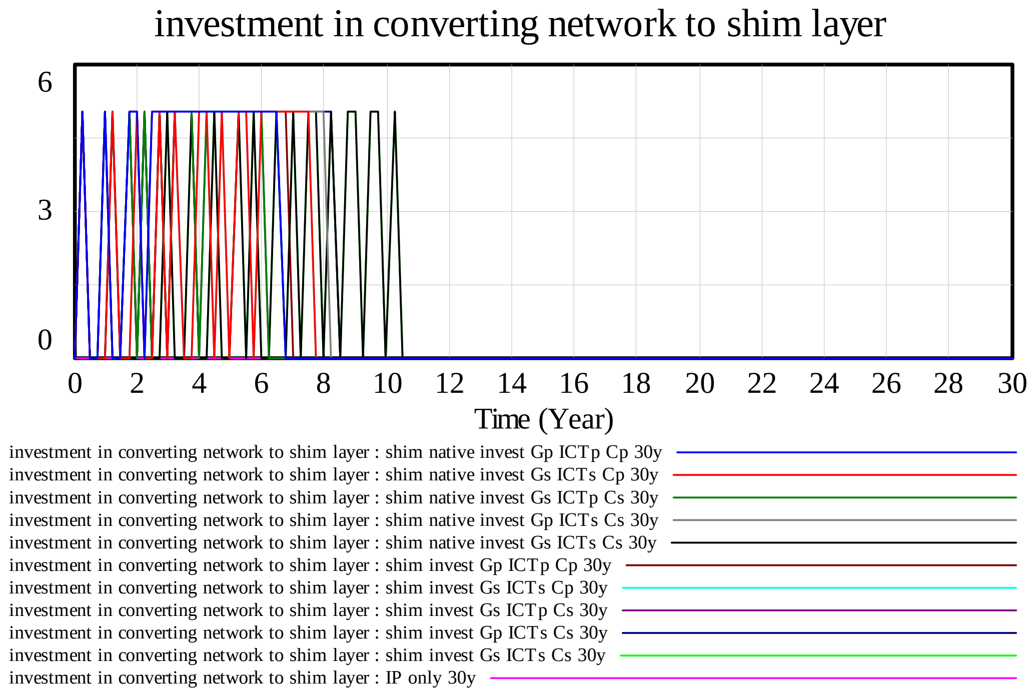

Figure 7.42: Network investment in enabling a 'shim layer' Information Cloud

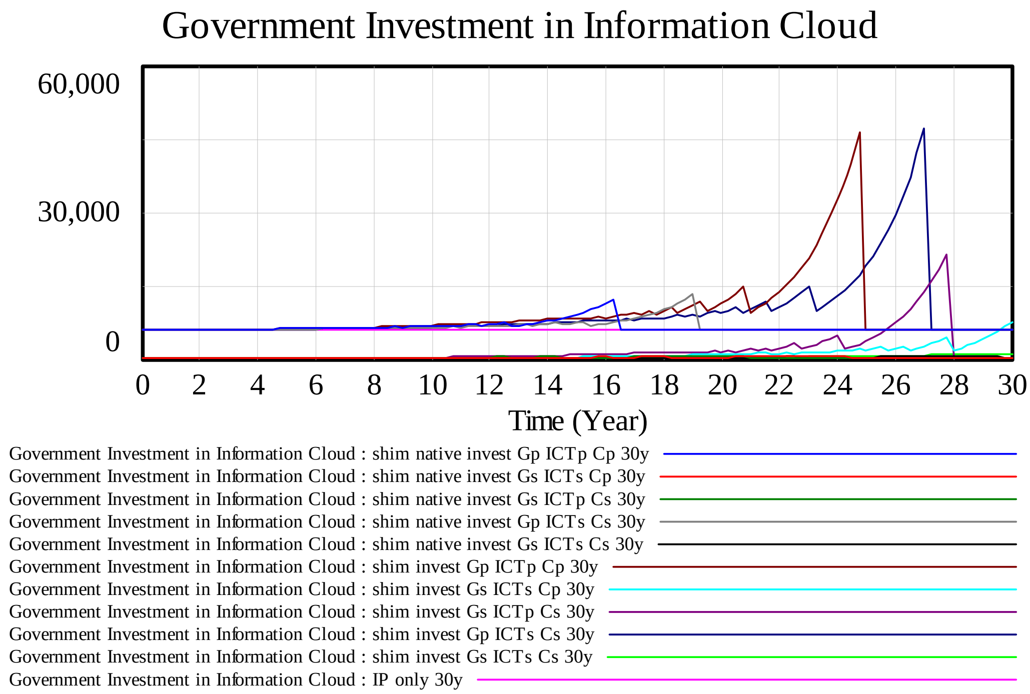

Figure 7.43: Government sector investment in Information Cloud applications over 30 years

In Figure 7.43 we see that the sub-scenarios in which the Government investment is proactive, but the network is reluctant to invest in the 'native' Information Cloud, result in the heaviest levels of Government sector investment. This is arguably an artefact of the model which links Government investment to 'shim layer' traffic levels. However, it is arguable that within a less well-developed Information Cloud infrastructure, investments from the public sector would need to be greater.

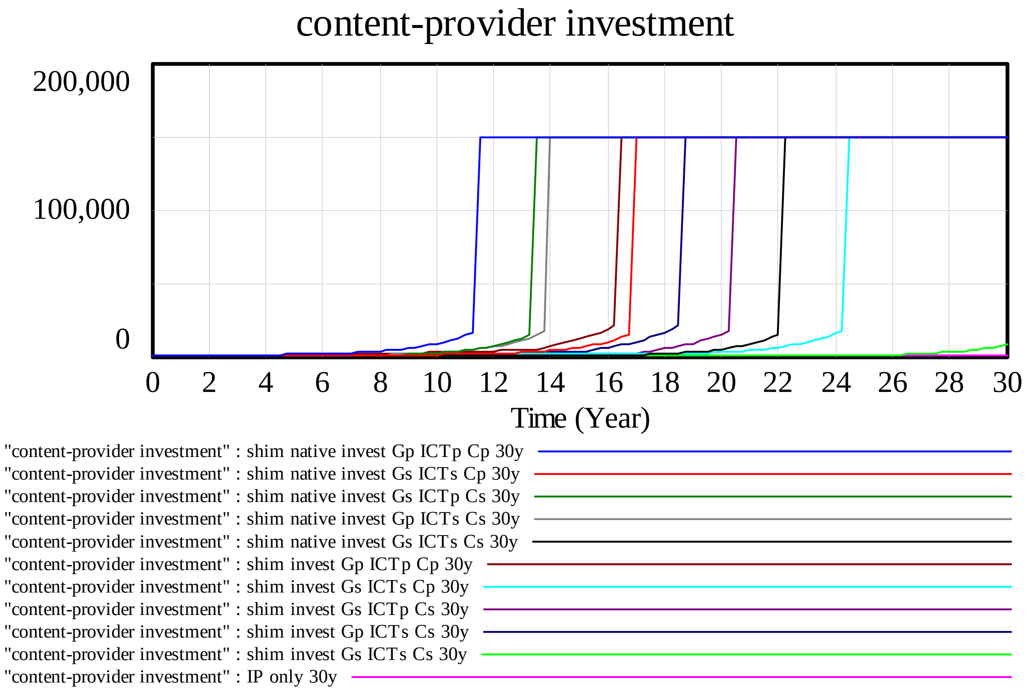

Figure 7.44: Content-Centric sector investment in Information Cloud applications over 30 years

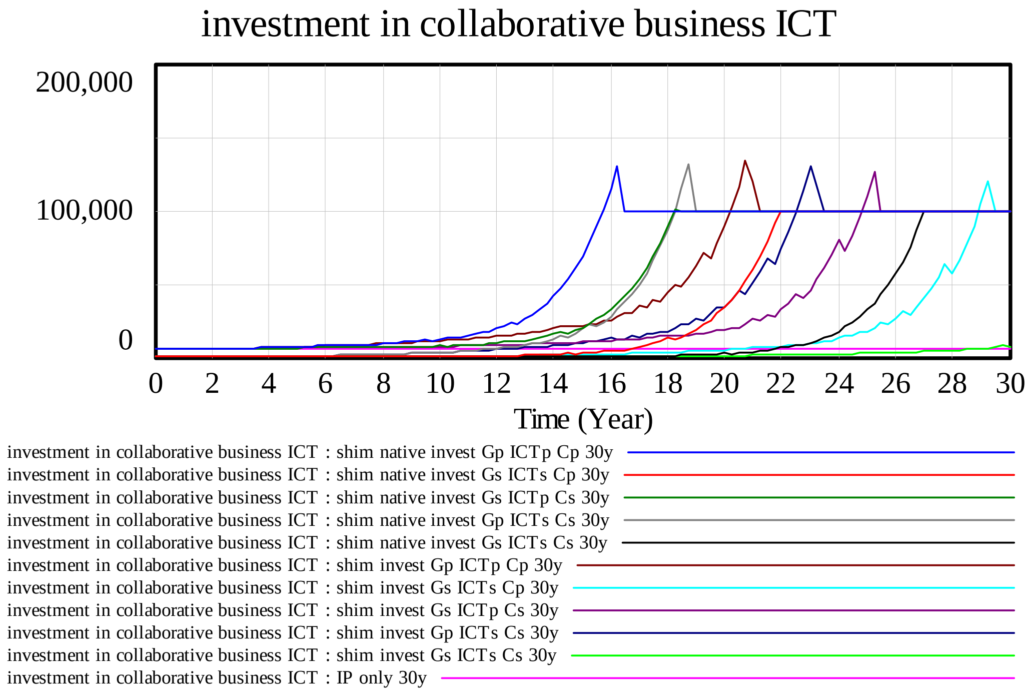

Figure 7.45: Collaborative Business ICT sector investment in Information Cloud applications over 30 years

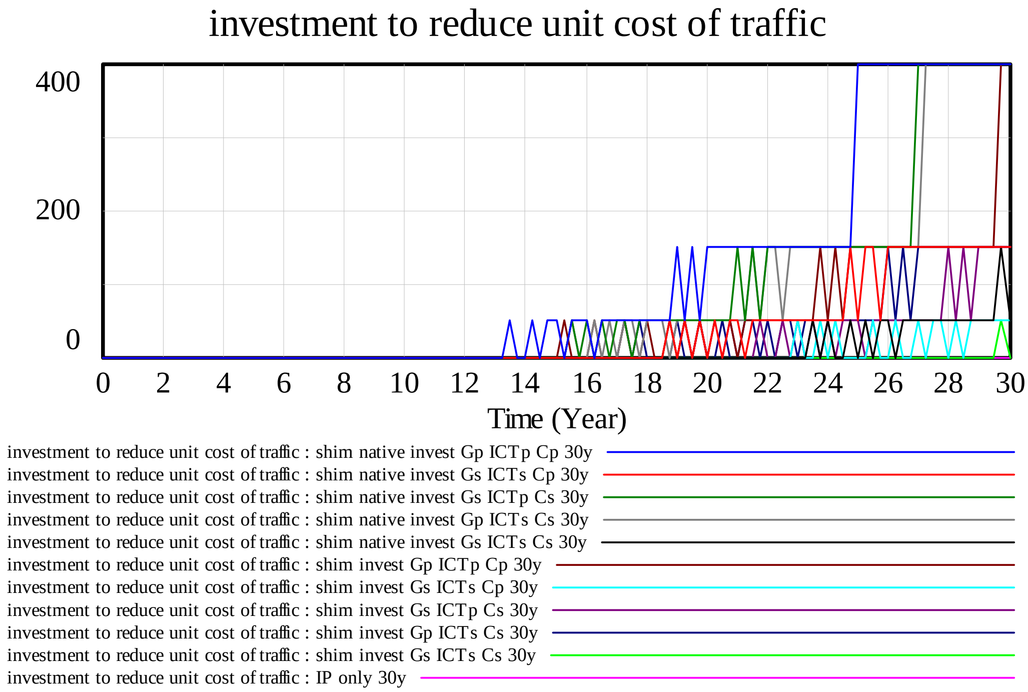

Figure 7.46: Network investment to reduce unit cost of traffic, over 30 years

8. Exploitation Analysis

Following on from the work defining scenarios and their effects on Information Cloud uptake, this section considers what implications these scenario outcomes provide, mainly from the perspective of an incumbent telecommunications operator. For example, which sector is critical to be convinced of the benefits of the Information Cloud? Or is it the case that without all three sectors being invested in, the uptake of the Information Cloud will stall? To what extent can different sectors operate in isolation, with only their own virtual private clouds? To what extent must the network owner be proactive?

8.1 Importance of Network and Sector Investment

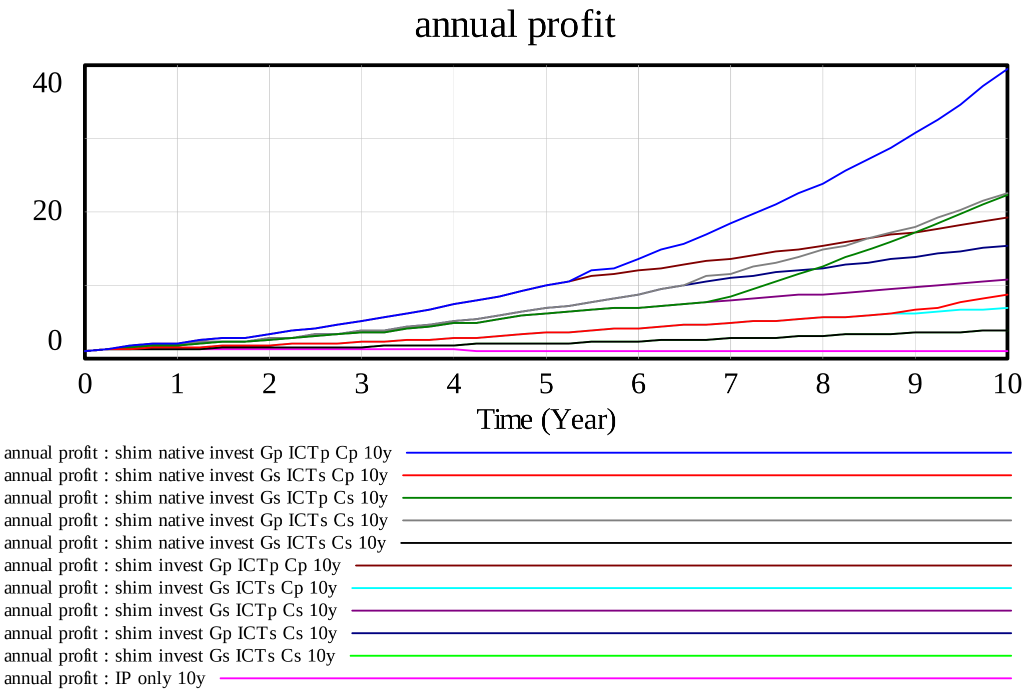

Figure 8.1: Reprise of network operator annual profit for the sub-scenarios over 10 years

In Figure 8.1, we review the modelled annual profit levels arising over ten years, for the different investment scenarios. Note that annual profit levels are before network investment costs have been deducted. Here we see that the most profitable scenario is the one in which all three application sectors: Government, Collaborative ICT and Content-Centric are proactive in their investment, and the network invests early in both shim and native Information Cloud transformation. However, this profit line takes 5 years to diverge from the line in which all sectors are proactive but the network only invests in the shim layer, showing that, even without taking into account the cost of network investment, profits from the native Information Cloud take 5 years to appear.

Comparing the relative importance of each application sector, in terms of its investment, we see that, over this ten year time frame, the most important sector investment is from the Government, although the annual profit line with only proactive Collaborative ICT investment (dark green), converges on the grey line: 'shim native invest Gp ICTs Cs 10y' at around ten years. This marks the Government as being an important early adopter in the predominantly 'shim layer' phase of the network transformation. This comes out of the model's assumptions of the key Information Cloud functionalities for Government, which do not require a native solution, in order to operate.

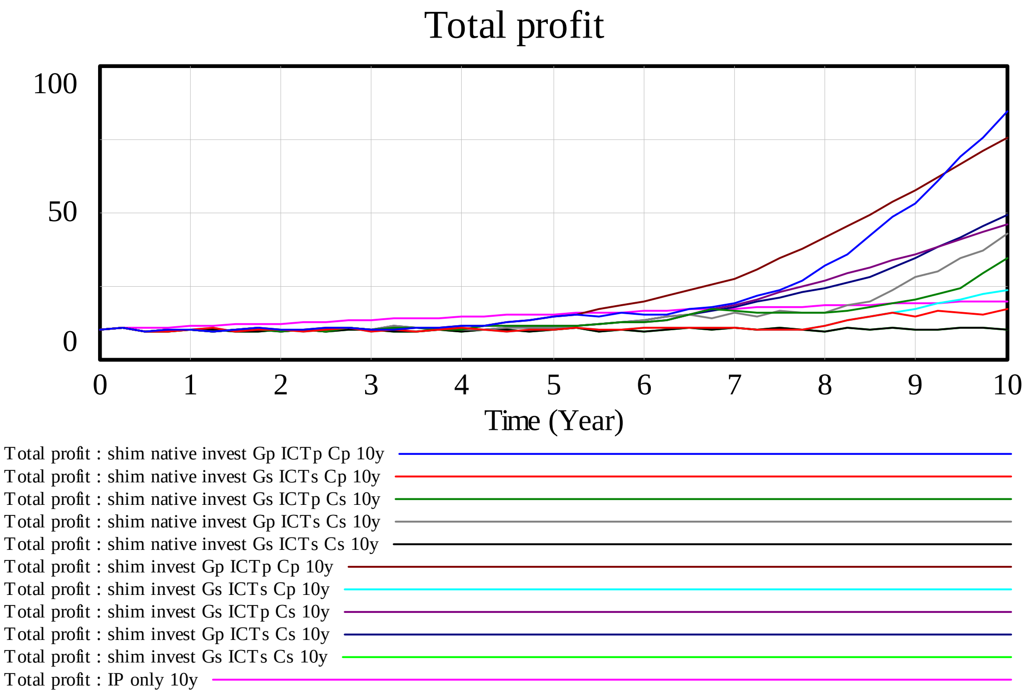

Figure 8.2: Reprise of cumulative total profit levels for the network operator for the sub-scenarios, over 10 years

Looking at Total Profit levels for the network operator, over ten years, in which its network investments have been taken into account (Figure 8.2) we see that it takes nearly ten years for the most proactive network investment 'shim native' to overtake the less proactive network investment 'shim', when all three sectors are proactive. So, even in the best-case scenario, proactive investment in the native Information Cloud has a less than ten year return curve.

If only one sector is going to be proactive over the first ten years, it would ideally be Government or Collaborative ICT. In this case, Total profit will remain slightly higher, over the ten year period, without proactive native Information Cloud investment. If the only sector that is investing proactively in Information Cloud is Content-Centric, then revenue growth will be very slow, with Total profit levels actually less than for the 'IP only' scenario (in which there is no network investment in Information Cloud') until around the ten year mark. Therefore, if there is no significant investment from Government or Collaborative ICT, network transformation will be a poor investment over a ten year timeframe.

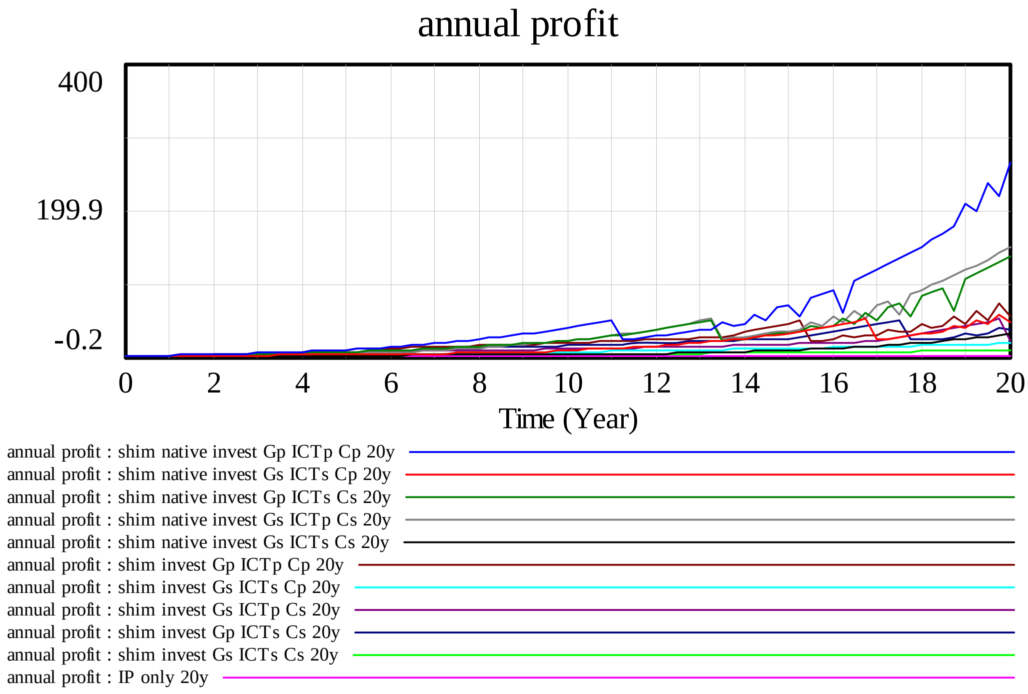

Figure 8.3: Reprise of network operator annual profits for the sub-scenarios, over 20 years

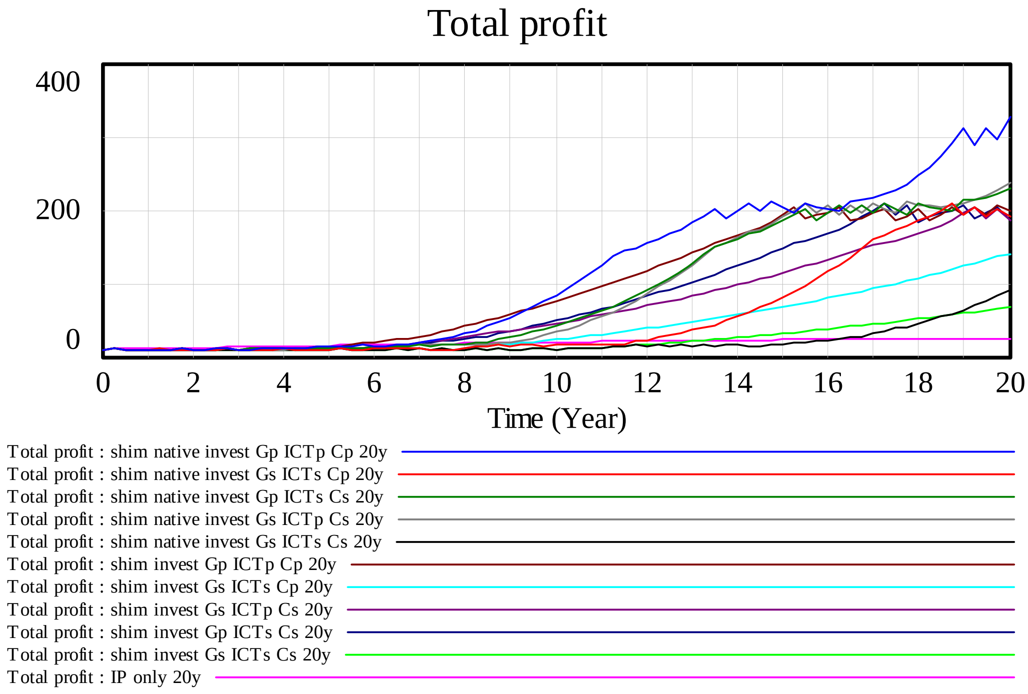

Figure 8.4: Reprise of cumulative total profit for the network operator, over 20 years

If we now look at annual profit levels over twenty years (Figures 8.3 and 8.4) we see that proactive network and sector investments make a huge difference to profit levels. Over 20 years, Collaborative ICT investment is slightly more important than Government investment.

Content-Centric investment remains a relatively weak factor in providing network profits. The scenario in which all three sectors are proactive, combined with sceptical native Information Cloud investment, is more successful than the scenario with proactive native Information Cloud investment, where only Content-Centric investment is significant. This is because investment in Content-Centric applications is strongly dependent on traffic volumes. So, if traffic is not being generated by other sectors, Content-Sector traffic will not grow rapidly either. This is based on the assumption that, unlike Collaborative ICT and Government, where Information Cloud solutions could be provided over a subset of the network, e.g. intra-Enterprise, Content-Centric applications rely on more general access to applications, i.e. in a much less regulated way, and so require other drivers of network transformation to be present.

Therefore, we envisage Government and Collaborative ICT as the engines for network transformation, with Content-Centric acting as a later adopter and a driver for the native Information Cloud, specifically.

Looking at the 'Total profit' curves, the results look less emphatic. This is because there is a modelled assumption that profits will be regulated, and that large traffic volumes will require additional investment in order to reduce unit cost. These effects act to level out performance in terms of 'Total profit', making the results appear less dependent on sector investment decisions. (A plateau in 'Total profit' levels is reached, for many of the scenarios, as further network investments are modelled, but the 'Total profit' curves do begin to rise again, after the plateau, as can be seen in Figure 8.4, looking at a 20 year timescale.)

If the model is correct, then one conclusion is that the worst performance, over twenty years, is for the 'IP only' scenario, in which there is no network transformation at all. However, in the absence of any significant application sector investment, the Total profits with either shim layer or native Information Cloud investment are worse for much of the twenty year period, only overtaking 'IP only' towards the end of the simulation. Application sector investment is therefore critical to the success of the network transformation, with even Content-Centric investment making a significant contribution to that success.

8.2 Final Analysis

In this section we summarise the conclusions of this exploitation analysis, underlining the key opportunity areas and the steps needed to exploit them.

In the previous section, we highlighted the importance of application sector investment to take place in concert with network investment, for network transformation to be successful. The conclusion of this is that the network operator has two main strategies:

[1] The network operator keeps a very close eye on developments and demand from application sectors, in order to time network investments;

[2] The network operator engages proactively with application sectors, and even provides some of the ICT solutions in order to catalyse the new products and services to run over the new network. Option 2) is likely to give the fastest network transformation, if done well. An incumbent operator who already has strong links and contracts with Government and business sectors and who also is capable to leverage network and ICT vendors, will be in the strongest position.

Expanding on option [2] which is the most proactive approach of the network operator, but also the most likely to succeed, we consider aspects of this strategy in more detail. According to the model, the application needs of the Government can be met from a 'shim layer' Information Cloud approach. That means that fundamental re-engineering of the network is not required. Here, the solution involves an overlay network with translation from information identifiers to an underlying IP network. Technologies such as Ericsson's MPSS (MPLS-like LIPSIN variant) could be used to achieve information-centric routing. Furthermore, many of the Government's issues could be addressed within a single tenanted solution. This is likely to be the right approach in transforming the Government's access to and organization of citizen data, for its own purposes. In dealing with this sensitive data, the Government would be unlikely to want to move straight to a 'true-cloud' multi-tenanted solution. However, in developing this solution for the Government, much would be learned about routing in an information-centric environment, and applications and services would be developed which could then be re-used in a more general context.

Continuing to look at the network operator working with its customers to develop information-centric solutions, we next consider the Collaborative Business ICT sector. This is likely to be more fragmented than Government, but will also involve potentially very large contracts. Many of this sector's needs can also be met by a 'shim layer' Information Cloud, but to achieve the full potential of Information as a Service and incremental business ICT growth, a multi-tenanted solution is likely to be needed. As such, solutions in the Collaborative Business ICT sector, would logically follow, or at least begin later than solutions designed for Government. Furthermore, to achieve the full set of functionalities for Collaborative Business ICT, such as dynamic bandwidth and infrastructure growth, and dynamic service composition, a 'native' Information Cloud, based more closely on a PSIRP architecture would be needed. However, initially this could be provided in a set of servers used for cloud-based services, and not require a complete Internet transformation.

For content-centric applications, there are two main drivers. The first is the driver for authentication of content, with the most commercially important requirement of this category being to have much tighter media rights management. This could involve information cloud authentication mechanisms to enable individual pieces of content to be licensed for specified periods to specified hardware or individuals only. We feel that although this would require mass-participation in order to be achieved, i.e. whatever mechanism were used, it would need to be widely accessible, it is quite a specific problem, that would require a specific solution, i.e. again it would not require a re-engineering of the Internet. However, it could be that it is the driver for Packet Level Authentication, which would benefit and underpin all other Information Cloud applications. As media rights management is commercially very significant, it would be a priority for a network operator interested in exploring the content-centric sector to identify the requirements for this sort of solution, and ideally, this would also have the effect of putting in place a key component of Information Cloud, in the form of PLA. The second main driver for content-centric applications is in achieving New Realities. Here we envisage many value-add products and services, such as 3-D gaming, haptic remote skill tuition, and immersive multimedia environments generally. Here, where dynamic bandwidth allocation becomes key, we see a strong dependency on a more fundamental re-engineering of the Internet than what can be achieved by using a simple overlay solution. Furthermore, in order for these services to be widely accessible, we reach a situation where a significant network transformation is needed (critical connective mass being needed) for these services to be commercially viable. We also see the dependence on a wide range of 3-D video and media mash-up technologies that may well fall outside of the scope of the network operator to be directly involved in. Therefore, we feel it likely that the content-centric sector will come to fruition later than the Government and Collaborative Business ICT sectors. Given that the services in the New Realities category are also more about lifestyle and would develop in response to fashion and trend, we would recommend that a more reactive rather than proactive response to this sector is followed.

9. Conclusions

Finally we reach our conclusions regarding the deployment incentives and business models for exploiting the PSIRP architecture, or, in other words, exploiting the Information Cloud.

In this work we have developed a methodology that takes the capabilities of a new architecture and identifies key application sectors that would benefit from those new capabilities. Following a combination of desk research and industrial liaison, stories were developed to highlight the business opportunities for each sector. These opportunities were then mapped to the 'Trigger points' or critical technological dependencies, in order to understand the specific requirements of the Information Cloud. In so doing, a distinction was drawn between Information Cloud capabilities that could be delivered using a 'shim layer' or tightly coupled information overlay; and those capabilities that would depend intrinsically on true or 'native' PSIRP being deployed over a re-engineered internet.

We developed a system dynamics model to link the investment decisions and demand for services being generated by the application sectors to traffic levels, network profits and investments. Although the parameter values were not based on known monetary values, the model enabled us to explore the impact of relative levels of investment. From this we were able to identify the inter-dependence of application sector and network operator investments. This led to the following main conclusions:

1. Application sector investment is critical to the commercial success of an Information Cloud, for the network operator.
2. Demand for the capabilities of Information Cloud exists and could be very attractive to the sectors explored: Government, Collaborative Business ICT and Content-Centric.
3. In order to optimise the timing of network investments and to stimulate demand, a network operator would ideally work in the Information Cloud application market, working closely with customers from the key sectors, to build the market opportunities.
4. The logical order in terms of markets to target would be Government, then Collaborative Business ICT, then Content-Centric.
5. A shim layer or information overlay approach should be pursued, in the first instance, with lessons being learned being transferable to a later native Information Cloud deployment.
6. In parallel with the main sector collaborations in point 4, there should be an early focus from the network operator on developing media-rights management, probably using Packet-Level Authentication. This is expected to be a strong early driver for the Information Cloud and PLA forms an important under-pinning of the wider Information Cloud.
7. Working with Collaborative Business ICT would be the most likely scenario for developing 'native' Information Cloud, firstly in a single-tenanted way (or multi-tenanted but with heavy safeguards, i.e. not general internet); moving eventually to a mass-market native information cloud for multi-media immersive environments.

10. References

[Can2005] C. Candolin, “Securing military decision making in a network-centric environment,” doctoral dissertation, Helsinki University of Technology, Finland, 2005.

[COR2009] ftp://ftp.cordis.europa.eu/pub/fp7/ict/docs/netmedia/20090507-position-paper-future-internet-architecture-with-disclaimer_en.pdf

[Cor2009] http://www.coriolisresearch.com/pdfs/coriolis_tesco_study_in_excellence.pdf

[FCN2009] Future Media and 3D Internet Task Force, *Future Internet and NGN: Design Requirements and Principles for a Future MEdia and 3D Internet,* Networked Media Unit, February 2009

[For2008] J. Forsten, K. Järvinen, and J. Skyttä, “Packet level authentication: Hardware subtask final report,” technical report [online], 2008, available at: http://www.tcs.hut.fi/Software/PLA/new/doc/PLA_HW_final_report.pdf [Accessed July 2008].

[Gar2009] Poor Security "The biggest culprits of data breaches in recent times - HM Revenue and Customes, Ministry of Justice, Department of Health and the Ministry of Defence - all reported a lack of basic systems for managing personal data effectively." Garlik, Government data policies continue to fail Britons, 2009

[IVI2009] "eCommerce 2.0 : The Multi-channel shift ", IVIS Group Whitepapers, 2009

[Jar2007] K. Järvinen, J. Forsten, and J. Skyttä, “FPGA design of self-certified signature verification on Koblitz curves,” Proc. of the Workshop on Cryptographic Hardware and Embedded Systems, CHES 2007, Vienna, Austria, September, 2007, pp. 256-271, Springer-Verlag LNCS 4727.

[Kob1987] N. Kobliz, “Elliptic Curve Cryptosystems,” Mathematics of computation, vol. 48, pp.203-209, 1987

[Lag2008] D. Lagutin, “Redesigning Internet - The packet level authentication architecture,” licentiate's thesis, Helsinki University of Technology, Finland, June 2008.

[MAR2009] Angelica Mari: http://www.computing.co.uk/computing/news/2241281/friends-provident-banks-tesco

[Max2009] L.Maxwell: http://www.cps.org.uk/cps_catalog/it's%20ours.pdf

[Mil1985] V. Miller, “Use of elliptic curves in cryptography,” Proc. of the Advances of Cryptology – Crypto '85, Santa Barbara, USA, August 1985.

[PSI2008a] S. Tarkoma (ed.), "PSIRP Deliverable 2.2.: Conceptual Architecture of PSIRP: including subcomponents descriptions", available at http://www.psirp.org/ [Accessed on 15 January, 2010].

[PSI2008b] M. Ain (ed.), “PSIRP Deliverable 2.3.: Architecture Definition, Component Descriptions, and Requirements”, available at http://www.psirp.org/ [accessed on 15 January 2010].

[PSI2010] M. Ain (ed.), “PSIRP Deliverable 5.5.: Final Plan for Using and Disseminating Knowledge”, available at http://www.psirp.org/ [accessed on 28 April 2010].