



Model documentation for the
Australian Competition and
Consumer Commission

Fixed LRIC model
documentation –
Version 2.0

August 2009
9995-204

Analysys Consulting Limited
St Giles Court, 24 Castle Street
Cambridge, CB3 0AJ, UK
Tel: +44 (0)1223 460600
Fax: +44 (0)1223 460866
consulting@analysys.com
www.analysys.com

Contents

1	Glossary	1
2	Introduction to model documentation	3
2.1	Purpose of this report	4
2.2	Key features of the model	4
2.3	Structure of the report	9
3	Demand projections	10
3.1	Services modelled	10
3.2	Data sources	14
3.3	Projections	14
3.4	Distribution of services across ESAs	17
4	Geographical analysis	18
4.1	Geotyping	20
4.2	Location analysis	27
4.3	Sampling	37
5	Architecture for the Customer Access Network	38
5.1	Overview	38
5.2	Basic algorithms	42
5.3	Traditional urban deployments	46
5.4	Traditional rural deployments	59
6	Customer Access Network Design	70
6.1	Services	70
6.2	Treatment of access technologies	71
6.3	Calculation of asset volumes in the CAN	72
6.4	Next-generation access network	79
7	Core network design – modern network	80
7.1	Demand drivers	83
7.2	Routeing factors (<i>Sheet Dem.Calc, Rows 91-270</i>)	86
7.3	Quality of service (<i>Sheet NwDes.1.Access, H10533:115792</i>)	86
7.4	Transmission modelling	87
7.5	Dimensioning at the LE level	88
7.6	LE transmission architecture	89
7.7	Dimensioning at the point-of-confluence (aggregation node) level	90
7.8	Dimensioning at the LAS level	92
7.9	Dimensioning at the TNS level	99
7.10	Special island solutions	105

7.11	Evaluation of access and core network overlap	106
8	Core network design – next-generation network	116
8.1	Demand drivers	118
8.2	Routeing factors (<i>Sheet DemCalc, Rows 273-375</i>)	118
8.3	NGN transmission	119
8.4	Access Tier 2 (<i>Sheet NwDes.1.Access</i>)	119
8.5	Access Tier 1 (<i>Sheet NwDes.1.Access</i>)	120
8.6	Aggregation nodes (PoCs) (<i>Sheet NwDes.2.PoC</i>)	120
8.7	Regional nodes	121
8.8	Core nodes	122
9	Costing calculations	124
9.1	Scenario control worksheet	124
9.2	Unit costs	124
9.3	Asset lives	127
9.4	Weighted average cost of capital	127
9.5	Tilted annuity calculation	127
9.6	Cost allocation between platforms	128
9.7	Treatment of common and indirect costs	130
9.8	Analysis of business overheads	131
9.9	Output results	143
	Annex A: Deriving the Location and Demand database	
	Annex B: Designing the sample of ESAs	
	Annex C: Clustering and spanning tree algorithms	
	Annex D: Fibre-ring deployment	
	Annex E: Assumptions in the access network deployment	
	Annex F: Cable network deployment	

© Commonwealth of Australia 2009. This report has been produced by Analysys Consulting Limited for the Australian Competition and Consumer Commission (ACCC). You may download material in the report for your personal non-commercial use only. You must not alter, reproduce, re-transmit, distribute, display or commercialise the material without written permission from the Director ACCC Publishing.

1 Glossary

<i>Acronym</i>	<i>Meaning</i>	<i>Acronym</i>	<i>Meaning</i>
ABS	Australian Bureau of Statistics	GoS	Grade of service
ACCC	Australian Competition and Consumer Commission	GRC	Gross replacement cost
ACMA	Australian Communications and Media Authority	HCA	Historical cost accounting
ADM	Add-drop multiplexer	HCPM	Hybrid cost proxy model
ADSL	Asymmetric digital subscriber line	HFC	Hybrid fibre-coax
AGW	Access gateway	ICCA	Interconnect call-charging area
ATM	Asynchronous transfer mode	IEN	Inter-exchange network
AUD	Australian dollar	IN	Intelligent network
BH	Busy hour	IP	Internet protocol
BHCA	Busy hour call attempts	IRIM	Integrated remote integrated multiplexer
BRAS	Broadband remote access server	LAS	Local access switch
BTS	Base transceiver station	LCS	Local carriage service
BUA	Built-up area	LE	Local exchange
CAGR	Compound annual growth rate	LPGS	Large pair gain system
CAN	Customer access network	LRIC	Long-run incremental cost
CBD	Central business district	LSS	Line sharing service
CMUX	Customer multiplexer	LTH	Local transmission hub
CPE	Customer premise equipment	MDF	Main distribution frame
DA	Distribution area	MEA	Modern equivalent asset
DDAS	Digital Data Access Service	MGW	Media gateway
DDAS	Digital data access service	MPLS	Multi-protocol label switching
DNS	Domain name server	MSAN	Multi-service access node
DP	Distribution point	MSPP	Multi-service provisioning platform
DSL	Digital subscriber line	MTH	Main transmission hub
DSLAM	DSL access multiplexer	NGA	Next-generation access
DWDM	Dense wave division multiplexer	NGN	Next-generation network
EPMU	Equi-proportional mark-up	NMS	Network management system
ESA	Exchange service area	NTP	Network termination point
FAC	Fully allocated cost	NU	Network unit
FDP	Final distribution point	OTA	Originating and terminating access
FLRIC	Fixed LRIC	PDH	Plesiochronous digital hierarchy
FTP	Fibre termination point	PIE	PSTN ingress/egress
FTTN	Fibre to the node	PoC	Point of confluence
GE	Gigabit Ethernet	Pol	Point of interconnect
GIS	Geographical information systems	POTS	Plain old telephony service
G-NAF	Geocoded national address file	PSTN	Public switched telephony network

<i>Acronym</i>	<i>Meaning</i>	<i>Acronym</i>	<i>Meaning</i>
PV	Present value	TGW	Trunk Gateway
QoS	Quality of service	TNS	Transit network switch
RADIUS	Remote authentication dial-in user service	TSLRIC	Total service LRIC
RAU	Remote access unit	TSP	Travelling salesman problem
RCU	Remote concentrating unit	ULLS	Unconditioned local loop service
RKR	Record keeping rule	UPS	Uninterruptible power supply
SDH	Synchronous digital hierarchy	USO	Universal service obligation
SEP	Signalling end-points	VB	Visual Basic
SIO	Services in operation	VDSL(2)	Very high speed digital subscriber line (2)
SIP	Session initiation protocol	VoIP	Voice over Internet protocol
SS7	Signalling system number 7	WACC	Weighted average cost of capital
STM	Synchronous Transport Module	WDM	Wave division multiplexer
STP	Signalling transfer point	WLR	Wholesale line rental
TDM	Time division multiplexer	xDSL	Generic term for DSL

Figure 1.1: *Acronyms used throughout the document [Source: Analysys]*

2 Introduction to model documentation

Analysys has been commissioned by the Australian Competition and Consumer Commission (ACCC) to develop a cost model of the modern and possible next-generation fixed network in Australia. This cost model (the 'Analysys cost model') estimates the long-run incremental cost (LRIC) of existing declared services for the period 2007–2012.

The model reflects Analysys's interpretation of the network deployment of a forward-looking, modern, efficient fixed network operator in Australia. It may differ in parts from the actual network deployed as the actual network has been affected by historical changes in the location and level of demand in Australia, as well as potential inefficient deployment strategies when viewed in a retrospective manner.

In addition to this document, reference should be made to:

<i>Fixed LRIC model user guide</i>	Provides guidance on use of the model, including the procedures for installing and running the model, and the location of the inputs in the active and offline modules of the model.
<i>Instructions for key processes in geoanalysis for the fixed LRIC model</i>	Provides an additional set of instructions for completing key data processing for the purposes of geoanalysis for the access network calculations in the fixed LRIC model.
<i>Description of the Visual Basic code in the fixed LRIC model</i>	<p>Outlines the principles and structure underlying the Visual Basic code developed for:</p> <ul style="list-style-type: none"> • the access network deployment • the core backhaul routes from local exchanges (LEs) to the local access switch (LAS) (or their next generation network (NGN) equivalents).

This model is highly complex and therefore it is recommended that this document and the user guide are viewed thoroughly in conjunction with the LRIC model. For parties interested in examining the Visual Basic code implementation, the code itself has been annotated should users want to inspect it within the Visual Basic Editor.

The default values in the LRIC model are benchmark values and do not necessarily reflect the ACCC's position.

2.1 Purpose of this report

The model documentation describes:

- the conceptual approach taken in the model
- the methodology applied in the model
- the assumptions made when developing the network design algorithms, where applicable.

2.2 Key features of the model

A principal aim of the model is to cost the specified declared services provided through the existing network, which in one form have been interpreted as delivering time division multiplexer (TDM) voice¹ over the fixed network using modern equivalent assets. However, the model can analyse a number of different scenarios associated with technology migration – ranging from producing a cost result for a fully modern network, through to producing a result for a fully NGN network. The set of these scenarios is controlled on a geotype level, where a geotype consists of a collection of areas which share the same common characteristics. Each geotype can be set to require either NGN network equipment or modern TDM network equipment.

The NGN approach modelled is based on a fibre-to-the-node (FTTN) architecture, where multi-service access nodes (MSANs) are deployed closer to the end user, instead of at the local exchange, and use the shorter distribution network copper cables to connect to end-users.

The modelling of NGN assets was a requirement of the project specification, and can facilitate analysis of how the steady-state cost of a service may change under a significant alteration in network architecture. It is not intended to have any relevance to ongoing discussions around NGN broadband access in Australia.

The model uses a parameterised approach which enables scenario testing to be performed on each of the input parameters. However, it is relevant to note that some modelling activities, including processing large exchange service areas (ESAs) through the geoanalysis and access network module, limit *rapid* testing of the end-to-end system.

The model allows pricing on a geotype basis, but also cost services based on the four ULLS bands used in current regulatory processes. The model also has the flexibility to consider different access pricing approaches.

The structure of the long run incremental cost (LRIC) model, which has been built primarily in Microsoft Excel, is outlined below:

¹ TDM voice is used to describe the modern network delivering time-division multiplexed voice services over a conventional PSTN network. This differentiates from voice-over-IP.

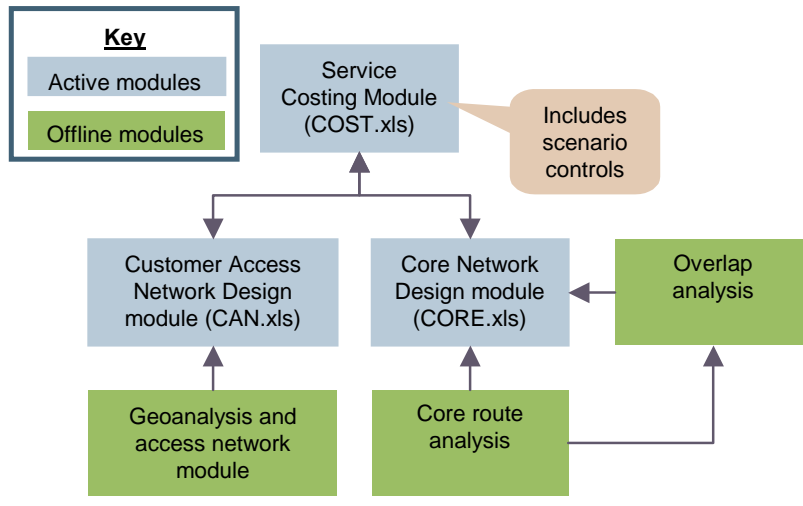


Figure 2.1:
Workbook structure of the
LRIC model [Source:
Analysys]

The LRIC model has two network design modules – a Customer Access Network (‘CAN’) module (CAN.xls) and a ‘Core’ module (Core.xls) – the network physical boundary between which has been identified as the equipment side of the MDF in the local exchange.

A combined Service Costing (‘Cost’) module (Cost.xls) has been developed for the access and core network assets deployed by the model. The Cost module has been populated with a benchmark set of values for costs and price trends. The model has the flexibility to allow for a different set of unit costs. Incurred capital expenditures are annualised according to a tilted annuity formula. This tilted annuity formula is discussed in detail in section 9.5.

Scenarios are controlled in the Cost module. These scenarios control the inputs that feed into the CAN and Core modules. The network outputs of these modules are subsequently fed back into the Cost module. Projections for subscriber demand (access lines and traffic) are contained within the Cost module.

The design of the access network is generated using parameters determined from detailed Excel and geographical information system (GIS) modelling work performed in external workbooks and database files, as described in section 4 and 5 and Annexes A–F. All of these files comprise the ‘geoanalysis and access network’ module.

The design of the routes within the core network – that which routes traffic from LE locations back to LAS locations via Point of Confluence (PoC) locations – is calculated in an external workbook (‘LE_LAS_ring.xls’). This is the ‘core route analysis’ module shown above in Figure 2.1 and is described in sections 7.6 and 7.7.

The ‘Overlap analysis’ considers realistic road routes for the core network routes and the overlap of these routes between the different layers of the core network and access network. This is described in section 7.11.

2.2.1 Key features of the Customer Access Network

The access modelling has taken a geotype approach with regard to its physical structure – Australia has been divided into a number of distinct geotypes. Network parameters have been determined using a sample set of the ESAs in Australia. Using these parameters, the information from the sampled ESAs has been extrapolated in order to size a customer access network for the whole of Australia.

The access network algorithms have been implemented in Visual Basic (VB) for Excel. In addition, detailed Australian geographical data, analysed using MapInfo (a GIS tool) has been used in order to define the efficient cabling structure and backhaul structures in the access and core layers of the modelled network.

A summary of the key features of the customer access network is provided below:

Sampled geotyping approach We have grouped similar ESAs into geotypes. This geotyping approach enables the extrapolation from a sample of ESAs up to the entire set of ESAs.

The geoanalysis and access network module completes a detailed calculation of the network assets required to serve a sample of over 800 000 locations within Australia. The asset volumes required for this sample are then scaled up in order to determine the asset volumes required for a full nationwide deployment.

Australian geoanalysis generating detailed access network calculations A dataset of locations derived from the Geocoded National Address File (G-NAF, sourced from MapInfo), have been mapped onto the street network, in order to help inform the network design process.

Inhomogeneous demand distribution methodology In the absence of data on the demand distribution, we have made modelling assumptions on the locations of businesses and the demand distribution across business locations (e.g. a few sites with a lot of lines, a larger number with fewer, etc.).

Street-aware cabling algorithms The access network algorithms take into account the location of streets in order to generate realistic network deployment architectures.

Where road networks cannot be directly traced, the model uses a calibrated adjustment to include the effects of street layout on inter-node distances.

Cost-minimizing transmission Detailed access network calculations are used which determine the efficient network layout from the customer to the local exchange using

<i>algorithms</i>	<p>cost-minimizing algorithms.</p> <p>Clustering algorithms and minimum spanning tree algorithms have been used to generate realistic trench and cable network architectures.</p>
<i>Transparent model results</i>	<p>For the sampled ESAs, maps of the access network can be generated; specifically:</p> <ul style="list-style-type: none"> • maps of the trench network deployed by ESA • maps of the nodes (distribution points, pillars, etc.) • maps of the fibre rings • maps identifying the clusters used.
<i>Model has the functionality to deploy four different types of access network solution</i>	<p>The access network algorithms deploy the cost-efficient solution for each modelled ESA. The algorithms deploy an appropriate network choosing between:</p> <ul style="list-style-type: none"> • copper • fibre • wireless (modelled using GSM 900MHz parameters) • satellite. <p>Cost proxy functions are used to generate realistic results in efficient levels of computational effort.</p> <p>In urban areas, the algorithms deploy either copper or fibre. In rural areas, the algorithms make a cost-based decision between copper and wireless solutions, and subsequently, between wireless and satellite solutions.</p>
<i>Flexibility in the tapering of copper deployment</i>	<p>A number of deployment approaches can be tested to support analysis of tapered versus non-tapered deployment assumptions. The default design choice is non-tapered and a summary of our own comparison is included in Annex F.</p>
<i>NGN access model</i>	<p>The access modelling framework allows for substitution of pillars by multi-service access nodes(MSANs) as part of an next-generation access solution.</p>

2.2.2 Key features of the core network

The Core module itself has been designed to generate results either for a modern network architecture – which is based on TDM transmission and SS7 signalling in the core; or for a next-generation network architecture – which is based on internet protocol (IP) transmission and session initiation protocol (SIP) signalling.

A summary of the model features is provided below:

<i>Australian-specific network architecture</i>	<p>The model has been specifically designed for the Australian situation:</p> <ul style="list-style-type: none"> • the modern core network architecture is based upon the existing network deployment • the scorched-node approach reflects existing location of local exchanges, local access switches and transit network switches • distances between nodes take into account the road/railway distance • the model takes into account special network solutions for islands in Australia.
<i>Ability to model modern and/or NGN network architectures</i>	<p>The core network model, includes autonomous calculations for the modern and next-generation network design architectures. The model can be run either as:</p> <ul style="list-style-type: none"> • deploying a modern (TDM and synchronous digital hierarchy (SDH)) network structure • deploying an NGN structure (SIP voice principally over Ethernet/IP-multi-protocol label switching (MPLS)), defined on a geotype basis.
<i>Node-by-node approach in the core</i>	<p>A node-by-node approach allows for each transmission link to be modelled individually – allowing specific equipment and transmission capacity thresholds to be modelled.</p>
<i>Efficient core network transmission that seeks to minimise cost whilst...</i>	<p>Specific transmission architectures have been modelled for each level in the core network:</p> <ul style="list-style-type: none"> • a clustering algorithm to identify PoCs (an LE on a transmission ring) and subtended small exchanges • a cost-minimisation algorithm to build rings through PoCs back to the LAS • ring structures have been modelled at the LAS-level • the transit links are modelled to be logically fully-meshed, but carried on a number of physical rings which seeks to reflect the existing network structure.
<i>... retaining a resilient network architecture</i>	<p>Where appropriate, resilient network architectures have been modelled to reflect the network design rules:</p> <ul style="list-style-type: none"> • where multiple-ring structures are required at the PoC level, rings are connected using two bridging nodes • LAS nodes are parented by two transit network switch (TNS) nodes • fully resilient ring structures are modelled at the PoC, LAS and TNS levels.

Transfer of costs between CAN & IEN Complex interactions between the CAN and inter-exchange network (IEN) are modelled and parameterised (allowing sensitivity testing), including:

- main distribution frame (MDF) use of the local exchange accommodation
- inter-exchange network (IEN) backhaul use of the CAN trench
- customer multiplexer (CMUX) and fibre backhaul use of the CAN trench.

2.3 Structure of the report

The remainder of this document is laid out as follows:

- Section 3 outlines the structure and assumptions made in the Demand module.
- Section 4 outlines the Australian geanalysis performed and the principles followed in developing the customer access network algorithms.
- Section 5 describes the architecture used for the customer access network design.
- Section 6 details the Excel calculation of the full customer access network assets.
- Section 7 defines the structure and assumptions made in developing the modern core network design algorithms.
- Section 8 defines the structure and assumptions made in developing the NGN core network design algorithms.
- Section 9 describes the costing calculations in the model.

The report also includes a number of annexes containing supplementary material:

- Annex A explains how the Location and Demand database has been derived.
- Annex B describes how the sample of ESAs has been designed.
- Annex C provides further details on the clustering and spanning tree algorithms.
- Annex D provides further details on fibre-ring deployment.
- Annex E outlines the assumptions made in the access network deployment.
- Annex F compares the difference in cost of a tapered and a non-tapered cable deployment for the sampled ESAs.

3 Demand projections

This section explains the demand projections which drive deployment of the access and core networks and allow recovery of costs over services. The following sections describe:

- services modelled
- data sources
- projections
- distribution of services across ESAs.

3.1 Services modelled

The project specification required that the following declared services² be included within the model:

- PSTN originating and terminating access services (PSTN OTA)
- unconditioned local loop service (ULLS)
- local carriage service (LCS)
- wholesale line rental (WLR)
- line sharing service (LSS).

However, to reflect the possible economies of scope, a number of other services are also modelled. Economies of scope, arising from the provision of both voice and data services across a single infrastructure, will result in a lower unit cost for both voice and data services.

Services modelled can be split between the two increments of access and traffic, and the following table provides a description of access services which drive the CAN deployment algorithms:

² Under Division 2 & 3 of the Part XIC of the Trade Practices Act, an access provider is obliged to supply a declared service to an access seeker upon request (Part XIC of the Trade Practices Act sets out a telecommunications industry-specific access regime).

<i>Service</i>	<i>Service explanation</i>
<i>PSTN end-user access</i>	Provision of a line suitable for voice and sold through Telstra retail. May be provided over copper pair, fibre (where multiple services may be provided over the same line), and wireless and satellite solutions.
<i>WLR</i>	Provision of a line suitable for voice and sold through operators other than Telstra retail. May be provided over copper pair, fibre (where multiple services may be provided over the same line), and wireless and satellite solutions. This service is not declared in band 1 but is still provided.
<i>ULLS</i>	The unconditioned local loop service allows an access seeker to provide services, including voice and/or xDSL, over the copper loop using its own equipment co-located with the termination point. The unconditioned copper pair runs between the NTP at the customer's premises and a customer access module (a concentrator at the MDF or a CMUX/MSAN deployed closer to the customer) – the copper pair terminates at a potential point of interconnection located at or associated with a CAM. The service includes the use of the equipment side of the MDF.
<i>ISDN BR access</i>	ISDN Basic Rate (BR) is a service comprising two voice or data channels (at 64kbps), which can be provided over a copper pair or fibre (where multiple services may be provided over the same line). As ISDN drives deployment of copper pairs and some backhaul requirements, we have included it in the service set.
<i>ISDN PR access</i>	ISDN Primary Rate (PR) is a service comprising up to 30 voice or data channels (at 64kbps), which can be provided over a copper pair or fibre (where multiple services may be provided over the same line). As ISDN drives deployment of copper pairs and some backhaul requirements, we have included it in the service set.
<i>Lines in the CAN</i>	In addition to the above services, other copper lines in the CAN which may be deployed can be quantified. This can affect the asset calculations in the CAN module and allows allocation of applicable costs away from the declared services.

Table 3.1: Services driving CAN deployment [Source: Analysys]

A number of voice services, including the PSTN OTA and LCS declared services, used in the model are listed in Table 3.2. These services contribute to the deployment of the core network (traffic increment).

<i>Service</i>	<i>Service explanation</i>
<i>PSTN local traffic</i>	<p>Voice calls between two Telstra retail subscribers within the same call charging zone. Estimates have been made to remove local Internet dial-up traffic from the projections.</p> <p>Calls are untimed so unit costs are presented on a per-call basis.</p>
<i>PSTN national long distance traffic</i>	<p>Voice calls between two Telstra retail subscribers not in the same call charging zone.</p> <p>Calls are timed so unit costs are presented on a per-minute basis.</p>
<i>PSTN outgoing traffic to international destinations</i>	<p>Voice calls from a Telstra retail subscriber to an international destination. The cost of the call is for the routing to Telstra's international gateway, but does not include the international costs.</p> <p>Calls are timed so unit costs are presented on a per-minute basis.</p>
<i>PSTN outgoing to mobile traffic</i>	<p>Voice calls from a Telstra retail subscriber to a domestic mobile operator. The cost of the call is for the routing to a point of interconnection (PoI), but does not include the interconnect cost charged by the receiving party.</p> <p>Calls are timed so unit costs are presented on a per-minute basis.</p>
<i>Dial-up Internet traffic</i>	<p>Circuit-switched calls made by customers for Internet access. Due to network routing and as these calls have a significantly longer connection time than voice, we have made estimates to remove this traffic from PSTN local traffic figures. Historical data and projections include traffic resold by access seekers.</p> <p>Under NGN scenarios, the dial-up traffic is reduced to reflect the number of subscribers connected to a MSAN where VDSL access is explicitly provisioned.</p> <p>Calls are untimed so unit costs are presented on a per-call basis.</p>
<i>PSTN OTA: PSTN originating/terminating traffic</i>	<p>A declared service which allows access seekers to sell voice services originating from a customer on Telstra's network or to terminate a call on Telstra's network.</p> <p>Domestic PSTN originating access is the carriage of telephone calls from the calling party (the A party) to a PoI with an access seeker's network. A PoI is usually located at a trunk (or transit) exchange.</p> <p>Domestic PSTN terminating access is the carriage of telephone calls from a PoI within an access seeker's network to the party receiving the call (the B party).</p> <p>Calls are timed so unit costs are presented on a per-minute basis.</p>
<i>LCS: Local carriage service</i>	<p>The local carriage service is a service for the carriage of telephone calls from customer equipment at an end user's premises to separately located customer equipment of an end user in the same standard zone. The service is used by access seekers to resell local calls.</p> <p>Calls are untimed so unit costs are presented on a per-call basis.</p>

Table 3.2: Voice services [Source: Analysys]

Services relating to Internet access which are included in the model are listed in Table 3.3. With the exception of LSS, the services are included to capture backhaul requirements from the LE towards the core network (traffic increment).

<i>Service</i>	<i>Service explanation</i>
<i>ADSL retail</i>	Provision of an asymmetric digital subscriber line (ADSL) Internet service and sold through Telstra retail. Currently, the cost of the line is recovered through the voice service. This service does not drive deployment of copper pairs, which are driven by PSTN and WLR services. Quantification of ADSL is important to assess backhaul requirements from an exchange and in the core network.
<i>ADSL wholesale</i>	Provision of an ADSL Internet service and resold by other operators. Currently, the cost of the line is recovered through the voice service. This service does not drive deployment of copper pairs, which are driven by PSTN & WLR services. Traffic from wholesale lines is handled in the same way as the ADSL retail service. No bitstream product is currently declared.
<i>SDSL retail</i>	As per ADSL retail, however the service is a symmetric Internet access product.
<i>SDSL wholesale</i>	As per ADSL wholesale, however the service is a symmetric Internet access product.
<i>LSS</i>	The line sharing service allows an access seeker to provide high-frequency services, principally xDSL, over the copper loop using its own equipment co-located with the termination point. Voice will be provided over low-frequency spectrum by Telstra.

Table 3.3: Internet access services [Source: Analysys]

In addition, the services in Table 3.4 are included in the model.

<i>Service</i>	<i>Service explanation</i>
<i>ISDN – voice traffic</i>	Associated with the ISDN access lines, voice calls from ISDN customers affect transmission requirements and share resources with the voice platform.
<i>Mbit/s in LE</i>	Transmission bandwidth between the LE and local transmission hub (LTH) is used by services identified above and other services (e.g. ATM, leased lines, connections to hybrid fibre-coax (HFC) and mobile networks, etc.) Projections of other services consuming bandwidth are made to reflect the economies of scope that would be reasonably achieved by the incumbent fixed-line operator.
<i>Mbit/s in LTH</i>	As above, but between LTH and main transmission hub (MTH).
<i>Mbit/s in MTH</i>	As above, but between MTHs.

Table 3.4: Other services [Source: Analysys]

Services modelled under NGN scenarios

An NGN can be partially or fully deployed, with the extent of deployment controlled at the geotype level. The services described above, do not explicitly change to NGN services as the geotypes are switched over to an NGN deployment, with the exception of dial-up internet traffic.

The principle aim of the LRIC model is to inform the ACCC and industry in pricing decisions for current declared services. It is unclear at this time what declared services may exist under new network architectures. Therefore the NGN version of the LRIC model can inform how certain voice and access service costs may change.

For access services, including ULLS and WLR, the principle driver of service cost will be the number of access lines over which the access network, albeit a modified network, is recovered over. Certain access services may not be applicable in an NGN scenario: this will be the case for ULLS due to the economics of deploying MSANs at each node. However, it is the total number of access service in use that impacts the unit costs of access.

For services on the core network, including OTA and LCS, the total volume of traffic may not change under an NGN scenario even if the voice protocol does. Dial-up internet traffic is decreased in proportion of the number of subscribers attached to a MSAN. The NGN scenario allows an understanding of how these existing declared services may change with network architecture. Under the NGN scenario, it is worth noting that the results will likely be sensitive to assumptions on the capacity provisioned for each subscriber for data services, which is a network input and discussed in section 8.1.

3.2 Data sources

Analysys has gathered historical data on the Australian telecommunications market from the following sources:

- company reports published publicly
- third-party estimates (Analysys Research, Globalcomms database)
- operator submissions to ACCC, under record keeping requirements (commercial-in-confidence)
- operator submissions to this specific project (commercial-in-confidence).

The data has been used to estimate the service demand in 2007 and to inform projections for the period 2008–2012.

3.3 Projections

The demand projections are contained within the Cost module and allow easy manipulation of demand projections, to support scenario testing.

Projections are based on simple compound annual growth rates (CAGR 2007–2012) and curves (see Figure 3.1) or the product market share in the long run. For more detailed information on the projected values, see the Inputs.Demand sheet of the Cost module.

Year modelled	2008	2011	2012	CAGR	Curve
PSTN access lines	YA	9,758,727	9,746,268		
Retail	YA	7,789,374	7,791,155		
Wholesale	YA	1,969,354	1,955,113		
WLR	YA	1,216,214	1,183,576		
ULLS	YA	753,139	771,537		
PSTN access lines	YE	9,746,268	9,746,268		
PSTN access lines growth	%	-0.3%	-%		
Retail	YE	7,785,295	7,797,014		
Wholesale	YE	1,960,972	1,949,254		
WLR	YE	1,197,600	1,169,552		
ULLS	YE	763,372	779,701		
% distribution of access lines					
Retail as % of PSTN access lines		80%	80%		
ULLS as % of Wholesale (WLR + ULL)		39%	40%		

Figure 3.1: Service projection in the Inputs.Demand sheet of the cost module [Source: Analysys]

CAN deployment

To dimension the CAN, a single projection is made which reflects the number of locations that need to be passed by the CAN. This is greater than the actual number of services in operation (SIOs) and reflects the fact that locations are passed which may not necessarily be connected in the year being modelled. The default value for this input is constant over time and set to be 11.3 million access SIOs. This is consistent with the Location and Demand Database that underpins the geoanalysis and access network module, as described in section 4.2.1. This default input can be changed by the user, if a different sized CAN is deemed appropriate.

The number of access SIOs over which the cost of the CAN is recovered, as a default input, is projected to decline from 10.4 million to 10.0 million SIOs between 2007 and 2012, as shown in Figure 3.2.

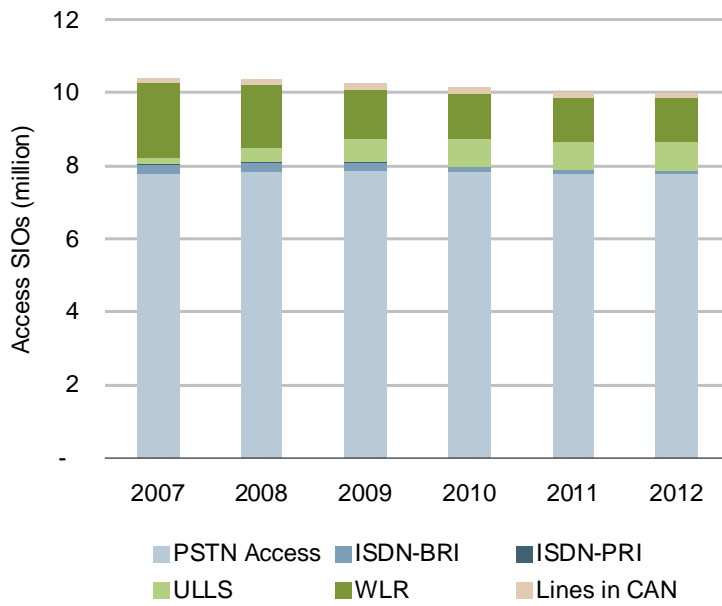


Figure 3.2:
PSTN-equivalent
access lines by type
[Source: Operator
data, Analysys
assumptions]

Voice services

Voice services estimates are based on historical usage per line (number of calls and average call duration), with assumptions made by service on year-on-year change. Figure 3.3 provides an illustrative projection for retail voice per retail line. LCS is modelled on a per-wholesale line (WLR) basis. PSTN termination and dial-up internet traffic is considered over the whole of the fixed access network PSTN lines.

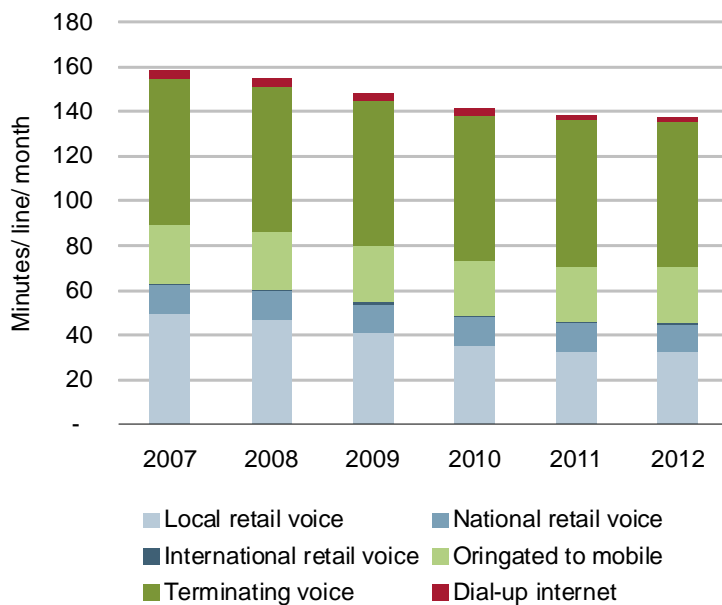


Figure 3.3:
Retail voice usage
per line [Source:
Analysys]

Internet services

Services related to Internet usage are based on assumed penetration of locations (household and business sites). Consideration is made for alternative access technologies including dial-up, HFC networks, wireless and satellite, leaving Internet services provided through xDSL (as illustrated in the figure below). Broadband connections are further sub-divided between retail, wholesale, ULLS and LSS services. Retail and wholesale xDSL services are further subdivided between ADSL & SDSL to reflect the platforms modelled in the core network.

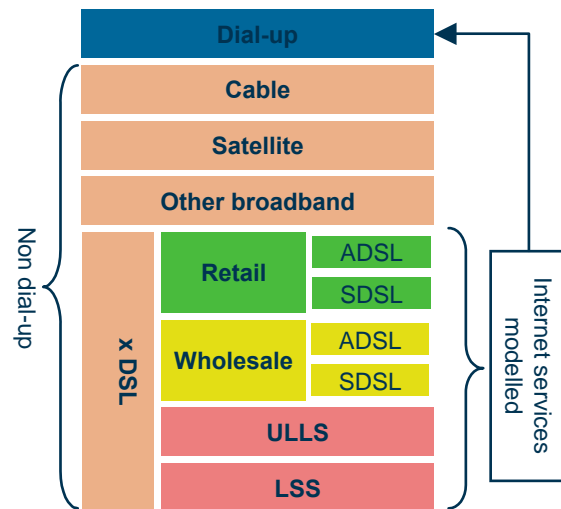


Figure 3.4:
Segmentation of
Internet services
market in Demand
module [Source:
Analysys]

3.4 Distribution of services across ESAs

It is necessary to estimate the distribution of access lines projections across ESAs for each year of analysis. The service line demand is distributed across ESAs on the 'In.Subs' worksheet in the Core module.

This worksheet distributes access lines based on the estimated demand database (see section 4.2), with locations reported by ESA. For PSTN, WLR and ISDN services (K33:N5286); for ADSL and SDSL services (AD33:AI5286); and for LSS services (S33:S5286) parameters are provided to toggle service availability by geo-type. It is appropriate to adjust services by geotype to reflect the availability of certain services (e.g. no copper access (ULLS) in Band 15 exchanges which are identified as satellite-only).

4 Geographical analysis

This chapter describes the geographical analysis (or “geoanalysis”) underpinning the access deployment algorithms, discussed in section 5. Australia is a large country with considerable diversity in its demography and topography, which needs to be understood and captured for any LRIC model of fixed services. To aid this understanding, we have:

- split Australia into suitable geotypes using a set of regions that cover the country
- undertaken “location analysis” to create a Location and Demand database
- generated a sample of regions in which to cost a network deployment.

The Geocoded National Address File was used as our primary source of locations, but this data set required some pre-processing since it was intended as a database of addresses rather than locations. For this and other reasons, it included far more entries than there were addresses in Australia. We undertook a filtering process to identify entries that were either repeated or superfluous to our modelling purposes. This is explained further in Annex A.

Having determined a subset of regions to be used as a sample, the volumes of network assets required to serve the demand within each region have been calculated using our network design algorithms, which are written using the Visual Basic compiler in Microsoft Excel. These algorithms have used the locations that lie within each of these regions, as well as their associated demand, as inputs.

These asset volumes derived by the algorithms from the sampled regions are then used to derive average volumes for each geotype. These average volumes are then applied to all the regions within a geotype, allowing the volumes required for an entire network of Australia to be derived.

MS Excel has been the basic tool in our geoanalysis work, but several more advanced calculations required a combination of the MapInfo, MapBasic and MS Access software packages. The processes that have used these packages are summarised in Figure 4.1 below.

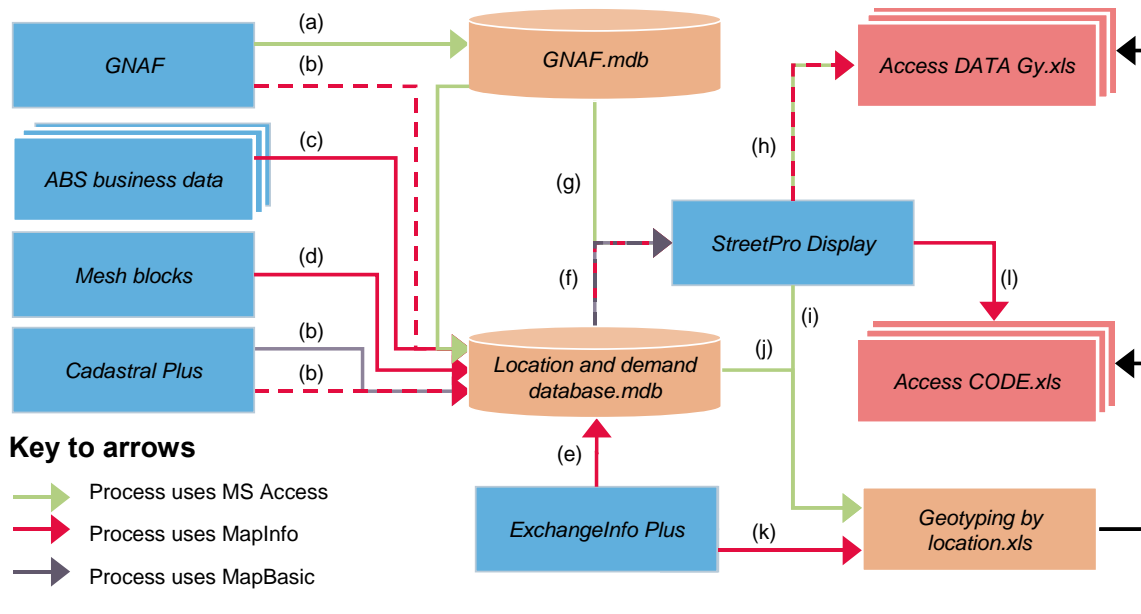


Figure 4.1: Summary of the geoanalysis workstream [Source: Analysys]

The data sets indicated above were obtained from MapInfo and the Australian Bureau of Statistics (ABS). The main calculations, labelled (a)→(l), are shown above.

Steps (a)→(d) pre-process the G–NAF into a database of addresses with associated demand:

- (a) The data contained within the G–NAF was pre-processed for use in our model. All 12.6 million entries were saved in MS Access and, as described in Annex A, several simple filters on the G–NAF fields were used to clean this database.
- (b) As an additional step in this filtering process, CadastralPlus together with the G–NAF were used in MapInfo to identify how the G–NAF entries were situated within the land parcels.
- (c) and (d) When assigning demand to the entries in the filtered G–NAF, MapInfo was used to designate each entry as residential/non-residential using the mesh blocks. MS Access was then used to assign a site size to G–NAF entries, guided by the mesh blocks and ABS data, as explained in Annex A.3.

Steps (e)→(h) aggregate the G–NAF entries into locations for the network deployment algorithms:

- (e) In order to identify the filtered G–NAF entries required for any sampled ESA, they were all associated with their parent ESA using ExchangeInfo in MapInfo.
- (f) and (g) When preparing the locations to be used by the deployment algorithms, the demand and co-ordinates of the filtered G–NAF were first imported from the MS Access databases into MapInfo and translated onto the street network of StreetPro. MapBasic was used to derive the co-ordinates of these translated points.

- (h) MS Access was used to group the G–NAF entries into locations and aggregate demand by location. MapInfo was used to translate each location into a suitable co-ordinate system.

Steps (i)→(k) produce outputs to inform the geotyping used in our model:

- (i) For the purposes of the geotyping, StreetPro was used in MapInfo to derive the total road in each ESA.
- (j) In addition, MS Access was used to aggregate G–NAF entries and locations by ESA, as well as derive the crow-flies distance of every entry from its nearest copper centre in its parent ESA.
- (k) Finally, ExchangeInfo was used to derive the area and number of copper centres within each ESA.

Step (l) is used to calibrate a “street-aware” distance function, called a p–function, with MapInfo processes used to break the street networks of the sampled ESAs into smaller segments.

This chapter and associated annexes highlight the calculations involved as follows:

- Section 4.1 describes how Australia was split into a series of geotypes
- Section 4.2 and Annex A outline the analysis of location data for the purposes of the geoanalysis and access network module
- Section 4.3 and Annex B describe how a sample has been generated for the geoanalysis and access network module.

4.1 Geotyping

4.1.1 Introduction

In order to reduce the calculations involved in the LRIC model, a geographical classification based on *geotypes* has been used to capture the diversity in Australian topography. Inputs to the network deployment on a geotype basis are then employed within the model.

Three things are required for geotyping:

- a set of regions that completely cover Australia, such as postcode areas or exchange areas
- geographical information associated to each region (e.g. population, area)

- understanding what model inputs would vary by geotype in the model or, at the very least, what factors drive these inputs.³

Geotyping aids the accuracy of using sampling, since average quantities derived from a sample for a single geotype should reflect the average quantities were they calculated across all regions in that geotype. Analysys has access to several geographical data sets for Australia shown below in Figure 4.2.

<i>Data set</i>	<i>Number of regions</i>
ULLS ESA classification (Telstra - public)	4
States (MapInfo)	8
Statistical Local Areas (MapInfo)	1329
ESAs (ExchangeInfo)	5070
Suburb and locality boundaries (CadastralPlus)	15 321
Distribution areas (ExchangeInfo Plus)	70 840
Mesh blocks (ACB)	293 980
Cadastres (CadastralPlus)	10 476 162

Figure 4.2: Australia data sets available to Analysys [Source: Analysys]

Classifications with either very few or very many regions are less suitable, since they (respectively) do not split out variation enough, or split it out so much that regions are indistinguishable from each other. On this basis, ESAs, suburbs and distribution areas (DAs) are the most suitable classifications for geotyping. The advantages and disadvantages are summarised in Figure 4.3.

³ For example, in mLRIC modelling, population density is seen as the fundamental factor underlying variations across geotypes. For our case, there may be several factors.

<i>Data set</i>	<i>Advantages</i>	<i>Disadvantages</i>
ESAs	ESAs would have been arranged to best support the population centres. Population centres have not significantly shifted in Australia, so these arrangements still have merit today	Tied to the existing network
	Telstra is able to provide data such as existing residential/business/xDSL services in operation (SIOs) at an ESA granularity. These will be helpful in the demand modelling	ESAs can still have population inhomogeneity, with a population centre surrounded by sparsely populated land
	The purpose of the model is to decide which technology(ies) are deployed within an area. ESAs have historically been used to make this decision from a copper perspective	
	Certain services (e.g. LCS) are not declared in certain ESAs. It would be preferable if geotypes did not have services declared in some parts but not in others	
Suburb areas	Suburb areas are more granular than ESAs	27% of the land (~0.3% of addresses) are contained within non-gazetted suburbs areas. This is a significant area with which to be using artificial constructs and may make the data set unreliable
	Suburb areas likely to be designed with some reference to population centres	Suburb areas do not lie neatly within ESAs, so existing data cannot be employed easily for calibration purposes
DAs	DAs are likely to be located at the population centres within ESAs	Not a complete covering of Australia
		Very heavily tied to the existing network The usefulness of DAs is uncertain given that, although the table implies a one-to-one mapping between DAs and ESAs, over a quarter of the DA objects lie across several ESAs. This may be a consequence of being added in for in-fill purposes

Figure 4.3: Advantages and disadvantages of each partition [Source: Analysys]

We have chosen ESAs because, not only are they a good guide to the population centres in Australia, but also they form the basis on which several regulatory decisions are defined, and the level at which data sets are reported which aid the calibration of the model.

4.1.2 Calculations

For each ESA, the quantities calculated for the geotyping process are stated below in Figure 4.4:

Quantity derived on an ESA basis	Source data
Pre-sifted G-NAF addresses	G-NAF
Pre-sifted G-NAF locations	G-NAF
Post-sifted G-NAF addresses	G-NAF
Post-sifted G-NAF locations	G-NAF
Adjusted, post-sifted G-NAF addresses	G-NAF
Adjusted, post-sifted G-NAF locations	G-NAF
Number of copper centres	ExchangeInfo
Area	ExchangeInfo
Total road	ExchangeInfo and StreetPro Display
Cumulative proportion of addresses in ESA within 1-10km (as the crow flies) of their nearest copper centre	ExchangeInfo and G-NAF
SIOs	Telstra
ULLS Band	ACCC

Figure 4.4: Quantities calculated [Source: *Analysys J Geotyping ESAs.xls*].

Some of these quantities, such as area and number of copper centres, are trivial to calculate. Others have required a more complex process. An example of this is the G-NAF database, which has been processed using the methods described in Annex A. Total road per ESA has been calculated by intersecting the street network in StreetPro Display with the ESAs in ExchangeInfo and calculating the length of the line segments lying within each ESA.

The cumulative proportion of addresses has been calculated using the co-ordinates provided in the G-NAF and the copper centre locations. For each ESA and every G-NAF entry in that ESA, the crow-flies distance is calculated to every copper centre in that ESA. The distance is calculated using a formula which accounts for the curvature in the Earth's surface, relevant when calculating distances across large ESAs. For two points with coordinates (LongA, LatA) and (LongB, LatB), the distance D between the two points is given by:

$$D = R * A \cos \left[\cos \left(\pi * \frac{(90 - LatA)}{180} \right) * \cos \left(\pi * \frac{(90 - LatB)}{180} \right) + \sin \left(\pi * \frac{(90 - LatA)}{180} \right) * \sin \left(\pi * \frac{(90 - LatB)}{180} \right) * \cos \left(\pi * \frac{LongA - LongB}{180} \right) \right]$$


where R is the radius of the Earth.

Certain ESAs contain multiple copper centres according to ExchangeInfo. For the purposes of calculating the cumulative proportion of addresses, a G-NAF entry is assumed to be served by the closest copper centre in the ESA, in terms of crows-flies distance. This simplifying assumption is reasonable since this calculation only plays an indirect role in the geotyping.

4.1.3 Methodology

The primary metric used for geotyping has been the *average road per location*, where locations are taken from our Location and Demand database. This is taken to be a proxy of the trench cost

required to deploy copper to the households. ESAs are grouped by both their ULLS band used in current regulatory practices and this ratio. Due to the variation in this metric (0.01–10 000m per location), intervals which double in size each time have been used, although in certain places these have been split out to analyse particular classes of ESAs with more granularity. This is particularly the case in Band 2, which contains a large number of high-demand ESAs. Figure 4.5 below shows the definition of the geotypes in terms of road per location.



Note: in this example, each colour represents a separate geotype (one colour for each of the geotypes).

Figure 4.5: Diagram of geotyping [Source: Analysys]

Bands 1 and 2

As shown above, this has worked well on Bands 1 and 2 and allowed the ESAs to be grouped into fairly narrow intervals, which are band-specific and as narrow as possible. Currently, Band 1 splits into two geotypes, whilst Band 2 splits into four geotypes. This minimises the variance of road per locations within these geotypes, increasing the likelihood that average values derived from a sample of ESAs will reflect the properties of the geotypes taken as a whole.

We must balance interval width against the number of geotypes, since this drives the sample size and in turn the running time of the network deployment algorithms.

Bands 3 and 4

Road per location as a geotyping metric is a good measure where a large proportion of the road will require trench along it. However, rural ESAs tend to have clusters of population connected by long roads that do not have trench deployed along it. In addition, the G-NAF is known to have address data of reduced accuracy in rural areas. Both of these effects lead to deceptively high

values of road per location. As a result, this ratio is less useful for the purposes of geotyping Bands 3 and 4 on its own.

Therefore, additional calculations have been included in order to refine the geotyping of these two bands, amounting to:

- defining each ESA in Bands 3 and 4 to be either clustered or spread
- isolating those ESAs with poor geographical data or having few/no services in operation (SIOs).

In addition, to minimise the number of geotypes, we consider the ESAs in Bands 3 and 4 together.

► *Clustered or spread ESAs*

An ESA is said to be *clustered* if more than a certain proportion of addresses from our filtered G-NAF database lie within a particular distance of their nearest copper centre in the ESA (if any). Otherwise, ESAs are said to be *spread*. A clustered ESA can be interpreted as some sort of significant settlement, usually containing a copper centre, surrounded by sparsely populated land. This effectively splits each of Bands 3 and 4 into two subgroups, giving four in total.

Bands 3 and 4 are classified into eight geotypes (7–14). Geotypes 7–13 contain either only clustered or only spread ESAs from Bands 3 and 4. Geotype 14 contains ESAs from Band 4 with very high ratios of trench per location. This geotype contains both clustered and spread ESAs, but contributes <1% of existing total SIOs.⁴

► *Isolating ESAs with poor geographical data*

An ESA is identified as having poor geographical data if it satisfies one of the following criteria:

- it contains no G-NAF entries
- no road has been identified in this ESA
- it contains fewer than a minimum number of SIOs, according to the data submitted by Telstra
- its ratios of SIOs per addresses and addresses per location both exceed certain limits: although urban areas are likely to have a high level of SIOs per address, for rural ESAs it is more likely that there are insufficient addresses currently included in the G-NAF.

Seventeen of the ESAs identified as having poor geographical data have been separately identified as either containing mines or resorts. Since these ESAs are likely to have demand concentrated at a small number of addresses, a high value of SIOs per address is reasonable in these cases. Therefore, these ESAs are allocated back into geotypes 1–14 using road per location as normal.

In total, over 500 ESAs are identified by these criteria as having poor geographical data. Since their location data are less likely to be a reasonable representation of where their demand is situated, we have decided to make these ESAs unavailable to the sample. However, their demand

⁴ Data is part of RKR requirements. Report is dated September 2007. SIOs include PSTN, ULLS and xDSL only.

still needs to be captured within the model, so they have been re-allocated to geotypes 7–14 using the ratio of road per SIO. We believe this to be a reasonable proxy of road per location, since the majority of rural locations would have one SIO per location. There are certain cases for which we believe neither ratio can be trusted and these are captured in two further geotypes (the 15th and 16th), as explained below. Of the ESAs with poor location data:

- Approximately seventy have fewer than 20 SIOs and are assumed to be served by satellite and placed in the 15th geotype.
- Seven have been identified as having no locations, no SIOs and no copper centre. These have been placed in the 16th geotype and have been excluded from the modelling altogether.
- Finally, two have been calculated to have no road, so both of their ratios are zero. These have been allocated to the 14th geotype.

4.1.4 Summary of geotypes

Figure 4.6 summarises the most important quantities associated with each geotype.

<i>Geotype</i>	<i>Band(s)</i>	<i>ESAs</i>	<i>Road (km)</i>	<i>Locations</i>	<i>SIOs</i>
1	Band 1	3	242	12 521	44 804
2	Band 1	14	502	8 977	237 009
3	Band 2	139	19 204	1 409 486	1 909 188
4	Band 2	219	37 942	2 117 145	2 825 812
5	Band 2	160	30 189	1 240 039	1 747 611
6	Band 2	66	12 004	319 370	494 792
7	Band 3/4 (clustered)	91	8 528	224 040	224 760
8	Band 3/4 (clustered)	107	19 150	132 038	153 554
9	Band 3/4 (clustered)	86	42 772	33 147	41 597
10	Band 3/4 (spread)	181	31 395	746 922	832 708
11	Band 3/4 (spread)	236	47 091	531 696	608 141
12	Band 3/4 (spread)	844	189 349	713 668	728 543
13	Band 3/4 (spread)	1936	540 340	399 316	395 470
14	Band 3/4	912	864 604	45 079	67 862
15 (satellite)	Satellite	69	29 094	508	695
16 (excluded)	No locations/SIOs	7	2 641	0	0
Total		5070	1 875 046	7 933 952	10 312 546

Figure 4.6: Summary of geotyped ESAs [Source: Analysys]

4.1.5 Current assumptions

The previous sections describe the various assumptions used by the geotyping calculations. The current values of these assumptions are given below in Figure 4.7.

<i>Assumption</i>	<i>Current value of assumption</i>
Threshold of SIOs below which an ESA is put into the separated geotype	20 SIOs
Maximum proportion of addresses that can be removed via sifting the G-NAF	50%
Determination of clustering in Bands 3/4: distance constraint	4km
Determination of clustering in Bands 3/4: minimum proportion of addresses in an ESA which must lie within the distance constraint of their respective copper centre	98%
Maximum ratio of SIOs per address in Bands 3/4, above which an ESA may be put into the separated geotype	2
Maximum ratio of addresses per location in Bands 3/4, above which an ESA may be put into the separated geotype	1.5

Figure 4.7: Current assumptions used in the geotyping
[Source: Analysys]

4.2 Location analysis

4.2.1 Location and Demand database

For a bottom-up LRIC model of this kind, ideally, one would have knowledge of the location of every access line required in Australia and the services demanded at each location. An access network could then be designed to serve these locations. Analysys does not have access to such a database, so has designed and built a reasonable proxy using available data.

The G-NAF provides basic geographical data for the location analysis, but this data has shortcomings for the requirements of the LRIC model. Although it provides a set of entries (or addresses) and corresponding locations, it does not provide any real indication of demand at these addresses. In order to complete a more detailed location analysis of Australia, several data sets have been used in addition to G-NAF (see Annex A for further information).

We have chosen to allocate high demand to addresses randomly within geographical areas likely to be served by high demand. The algorithm that we have designed has a control on:

- the number of addresses which are given additional demand
- the location of the addresses with higher demand
- the amount of demand allocated to these addresses.

We have attempted to calibrate the number of business sites of various sizes on a regional level. We have also used the mesh blocks to *inform* where business premises are likely to be and have focused our allocation of higher demand primarily, but not entirely, to these areas. Finally, by grouping G-NAF entries into locations, we have captured multi-dwelling buildings.

This leads to a set of G-NAF entries, each with an associated number of units of demand, that covers the whole of Australia. This data set is hereon referred to as the Location and Demand Database and is able to:

- (after grouping) provide a set of locations corresponding to actual points where demand for fixed services originate in an ESA and locations corresponding to be normal households, multiple dwellings and business locations.
- distinguish between business and residential addresses, to reflect the likelihood that businesses usually lie in a central area (or areas) and residential areas also tend to collect into one place.
- capture inhomogeneous demand, with normal households usually requiring one or two lines, business addresses multiple lines.

For any given sample of ESAs to be processed by the geoanalysis and access network module, a set of locations with associated demand can then be extracted.

The processes that we have used to derive this database are summarised below in Figure 4.8:

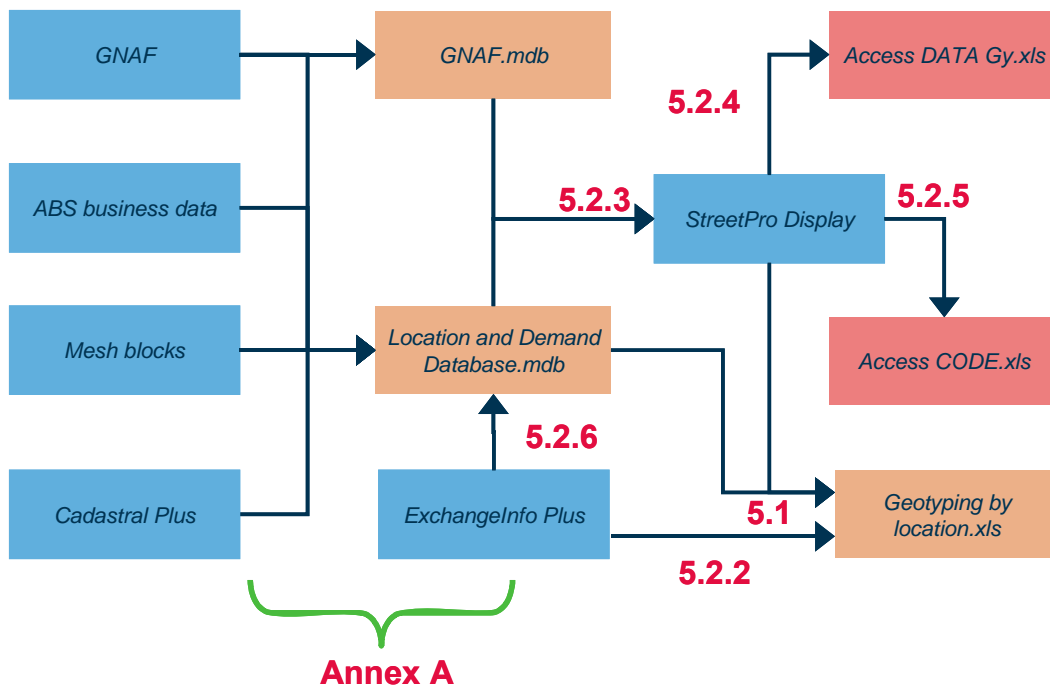


Figure 4.8: Location and Demand database and the corresponding chapters [Source: Analysys]

The above also indicates the geographical processing required in addition to producing the Location and Demand Database. We describe each of these processes briefly in the remainder of section 4.2. They cover the following topics:

- including the relevant copper centre for the area from ExchangeInfo
- mapping the locations onto the street network
- converting the coordinates into a suitable coordinate system rather than latitudes and longitudes to use in a distance function
- deriving a p-function to use a “street-aware” measure of distance
- the treatment of ESAs with multiple copper centres.

4.2.2 Copper centres

In order to satisfy the scorched-node assumption for the bottom-up LRIC model, the copper centres provided in ExchangeInfo are included as a site for a node in the region. These points are taken exactly as is from the data set.

4.2.3 Mapping locations onto streets

The Location and Demand database almost exclusively uses G-NAFs that have higher reliability, meaning that the location is somewhere within a land parcel. In reality, the location of the network termination point (NTP) is highly unlikely to coincide with this location within the land parcel provided by G-NAF.

Therefore, we mapped locations onto their nearest street to a point which we call the final drop point (FDP), as shown in Figure 4.9. We then calculate the trench and cable length required to reach the exchange from that final drop point (FDP). The distances associated with the actual final drop, from this FDP to the NTP are derived using separate assumptions (e.g. average street frontage, average street width etc.) in the CAN module.

We have used the Distance Calculator in MapInfo 8.0 to map our locations onto the set of lines in StreetPro representing the Australian road network. A simple MapBasic program has been written to identify the latitude and longitude of both ends of the lines that the Distance Calculator generates.

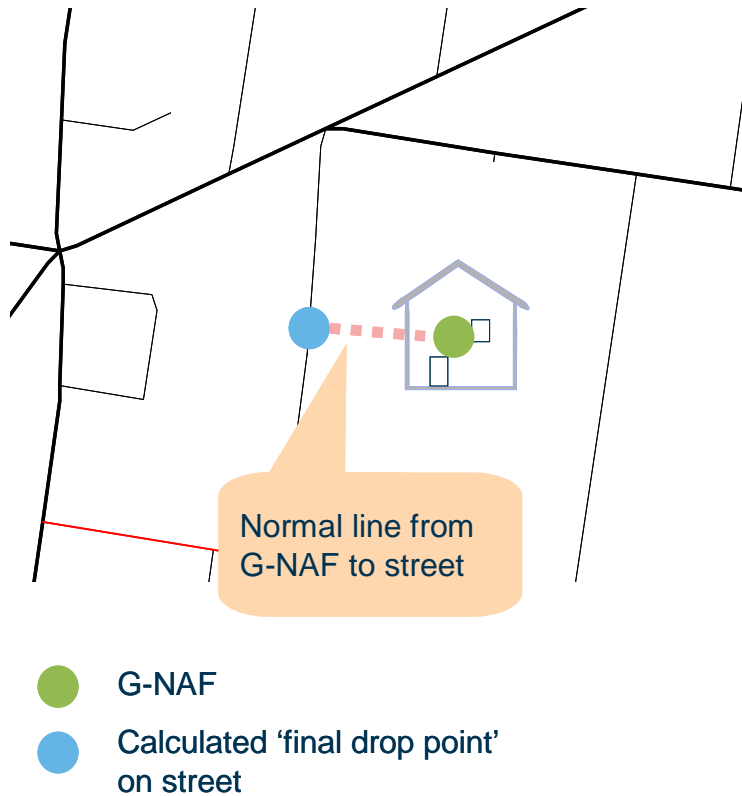


Figure 4.9:
Mapping G-NAFs onto
the streets using
StreetPro Display
[Source: Analysys]

In order to ensure that the points remain within each region, only lines representing street network within the ESA are used for this process. This is found by identifying which road segments intersect the region and then erasing any part of these segments that lie outside the ESA boundary. An example of the difference that this can make is shown for the ESA Exhibition in Figure 4.10.

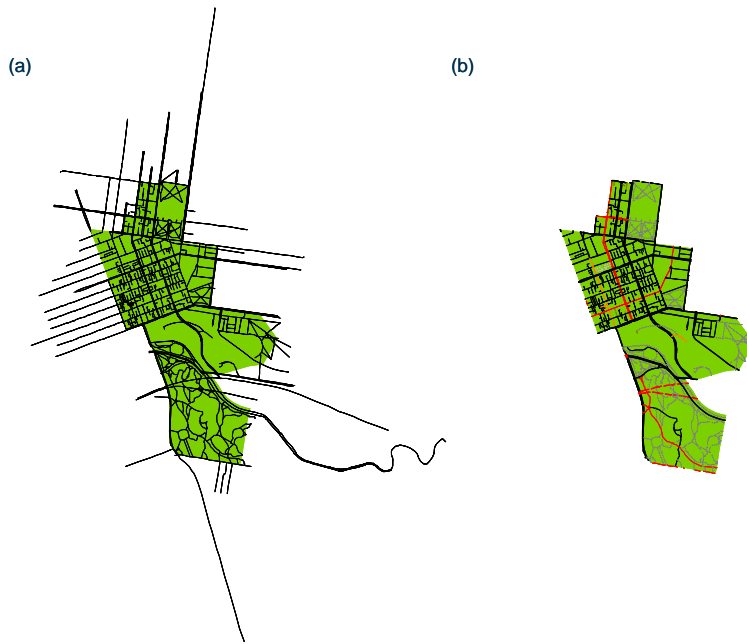


Figure 4.10:
Street line segments from StreetPro Display that (a) intersect the ESA and (b) actually lie within the ESA. We have used (b) to identify a suitable point on the street [Source: Analysys]

4.2.4 Coordinate system

In order to capture realistic street distances in the LRIC model for calculating trench and cable lengths, we are using a p–function to derive the distance between two points (see below). Latitudes and longitudes cannot be used directly with this function: a coordinate system that is equivalent to calculating in metres is required. Therefore, any coordinates used in the model are calculated with the Australian map grid (MGA94). These values can be calculated using the Coordinate Extractor tool in MapInfo. However, Australia is a large country and the exact coordinate system used depends on the longitude. The “zones” used in the Australian map grid are shown below in Figure 4.11.

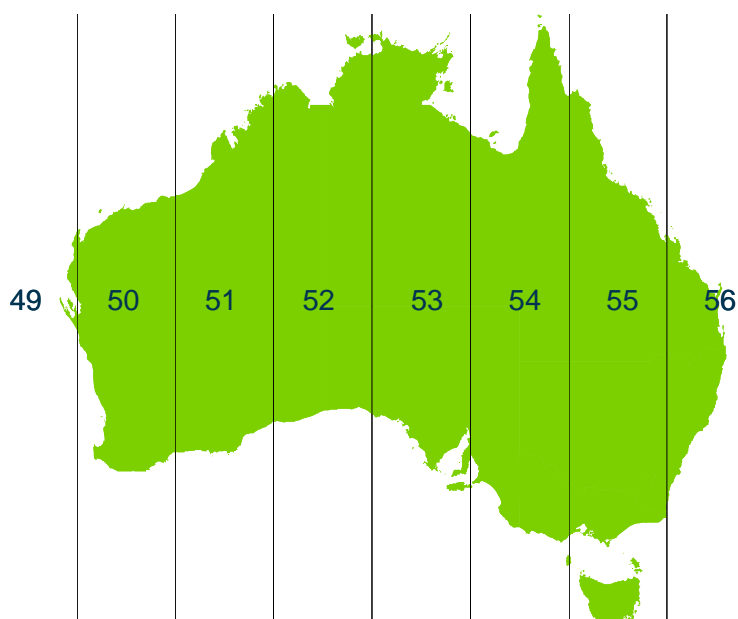


Figure 4.11:
Different zones in MGA94
[Source: Analysys]

When deciding on a coordinate system for a set of locations to be used by the geoanalysis and access network module, the rule that we have employed is that the zone containing the majority of the locations is used. This is because different coordinate systems cannot be mixed when calculating distances between points using the p–function.

4.2.5 Deriving p–functions

A p–function is a generalised form of measuring distance between two points. Crow-flies, or Euclidean, distance is the simplest measure of distance and is the length of the straight line joining the two points. However, our model needs to be capable of estimating the distance between two points but only by moving through a network of roads i.e. the model needs to be “street-aware.” We have therefore calibrated a more complex p–function to achieve this. This calibration involves:

- taking a large sample of point-to-point links across all sampled ESAs
- calculating the actual road distance of each of these links
- identifying p–functions that most closely derive these road distances for the links using just the co-ordinates of the link endpoints.

Definition

A general p–function is defined by two parameters k and p and defines the distance between two points $\underline{x}=(x_1,x_2)$ and $\underline{y}=(y_1,y_2)$ by:

$$\|\underline{x} - \underline{y}\|_{k,p} = k \left(|x_1 - x_2|^p + |y_1 - y_2|^p \right)^{\frac{1}{p}}$$

Euclidean ‘crow-flies’ distance uses k=1, p=2; whilst rectilinear ‘Manhattan’ distance uses k=p=1. At a local level, we have assumed that it is acceptable to use crow-flies to estimate distances. However, over larger distances, it is better to use a different pair of values to account for the fact that, in reality, trench and cable will follow roads and this will not be captured by either crow-flies or Manhattan distance alone. With regard to identifying suitable values for k and p, a paper by Love et al⁵ concluded that “it is inappropriate to simply assume a convenient value for p in a location study.” This is because experimental data finds that there are several values of p that can lead to good approximations. Also, different values of k can be used to scale distances to be a better approximation.

We believe that the treatment of the lengths of cable and trench need to be slightly different within the geoanalysis and access network module, as a single trench can actually hold multiple cables. An example of this is shown below.

⁵ Love, Morris and Wesolovsky,(1988) Facilities location models and Methods, Ed Saul J Goss, North Holland.




Figure 4.12:
 Demonstration of trench
 sharing between links
 [Source: Analysys]

Figure 4.12 above shows two links joining FDPs to their parent DP. Diagram (a) indicates how these links would be viewed links if trench and cable used the same distance measure. Diagram (b) shows what is happening between these links for our example in reality: they are sharing part of the total trench as the locations lie across a junction. Although a calibrated p -function would get the cable lengths correct, it would over-estimate the actual trench. Therefore, Analysys has calculated a separate trench sharing coefficient, j , in addition to determining values of k and p . Cable length is calculated using the p -function, whilst the total trench lengths are calculated using the p -function and then scaling by j .

Both the p -function and value of j can vary by geotype in the algorithms, although crow-flies distance is still used at certain times in the code, in particularly whenever comparing distances in the clustering algorithms.

Calculation of j , k and p

Using a p -function to calculate cable and trench distances by geotype requires the derivation of suitable j , k and p . We have calculated these parameters using the following steps:

- Disaggregate the StreetPro dataset into a dataset that can be used for plotting routes
- Generate a trial set of links by calculating the network for the sample using crow-flies distance
- Select a sample set of these trial links from each ESA
- Calculate the actual route for each sampled links, as well as the corresponding road distance
- Filter out the trial links that are unrealistic
- Calculate the parameters k and p by geotype
- Calculate j by geotype.

These steps are described in further detail below.

We have calculated values for j , k and p for each geotype using the street network from StreetPro lying within each of the sampled ESAs. However, this dataset contains fairly long lengths of road and these segments do not always terminate at junctions. As a result, we have used functionality within MapInfo to break the line segments into smaller lengths that do terminate at junctions.


We calibrate the p -function using reasonable links between real locations generated by our algorithms. To that end, the geoanalysis and access network module has been run for the full sample of ESAs using only 'crow-flies' distances. Although the distances generated by this calculation are not correct, the clusters and spanning trees generated are reasonably similar to

networks generated using a calibrated p–function. The algorithm outputs a set of crow-flies trench links, specified by their endpoints.

From the Euclidean-based outputs for each sampled ESA, a link identifying the incremental trench link between two nodes⁶ has been selected at random and extracted. In addition, the full set of trench links joining all the locations served by these two nodes have been extracted.

A software package called RouteFinder (version 3.41) has then been used to calculate the shortest road distance between the endpoints of these extracted links. This is completed using the software’s Route Pairs functionality.

Since the first stage of the network design algorithms, specifically the clustering algorithms, are not street-aware, some filtering of anomalies is required. For example, two points can be stated as being linked despite the absence of direct road between them. RouteFinder will then calculate exceptionally long links which would never be used in reality. An example is shown below in Figure 4.13.



In order to remove such anomalies, we have filtered our sample of links on the assumption that a link is an anomaly if the ratio of road distance to crow-flies distance exceeds 2.5. This sifted set of links have been grouped by their parent geotype.

For each geotype, Excel Solver is then used to identify the values of k and p that minimise the total sum of the quantity:

$$(\text{Road distance} - k*((x_1-x_2)^p + (y_1-y_2)^p)^{1/p})^W$$

where W is a constant. Several values of W were tested and it was found that $W=1.5$ generated p–functions that best approximated the total link length across all geotypes. This total sum is

⁶ Where the nodes can be any combination of (pillars, LPGS or RAU)

calculated across all links (defined by (x_1-x_2) and (y_1-y_2)) in the geotype, but excludes all links within DP clusters.

In order to calculate the trench sharing coefficient j for each geotype, the total non-overlapping distance n of all links is calculated (effectively the trench network). j is derived using the ratio $j = n / d$, where d is the total length of all individual links in the geotype derived using the distance functions. Specifically, links within DP clusters are calculated using Euclidean; all other links are calculated with the p -function using the values the values of k and p previously determined. j is then used to scale total trench length for all types of links. This includes not only links calculated using the p -function but also the links within DPs.

Values of k and p and j were calculated for all 14 geotypes, and are given below in Figure 4.14.

Geotype	k	p	j
1	1.16	1.34	0.82
2	1.33	2.30	0.88
3	1.04	1.21	0.90
4	1.03	1.26	0.93
5	1.23	2.30	0.97
6	1.14	1.47	0.87
7	1.20	1.63	0.84
8	1.12	1.56	0.85
9	0.99	1.20	0.94
10	1.25	1.87	0.85
11	1.06	1.32	0.86
12	1.16	1.63	0.87
13	1.17	1.38	0.88
14	1.23	1.90	0.89

Figure 4.14: Parameters for the p -function for each of the geotypes [Source: Analysys]

Values of p between 1 and 2 are intuitive, lying somewhere between Manhattan ($p=1$) and Euclidean ($p=2$) distance, as demonstrated below in Figure 4.15.




Figure 4.15:
Interpretation of p -
functions for $p=1$,
 $p=2$ and p between 1
and 2 [Source:
Analysys]

Two geotypes use a value of p much greater than 2. This is less intuitive, but still valid, as each such pair of k and p still always return distances that are longer than crow-flies.

4.2.6 ESAs with multiple copper centres

Over 100 ESAs have multiple copper centres and some of these appear in the sample. We have subdivided these ESAs individually using a Voronoi construction in MapInfo, based on the location of these copper centres within the ESA. An example is shown below in Figure 4.16.




Figure 4.16:
The Voronoi construction
for an ESA containing two
copper centres (the black
squares) [Source:
Analysys]

The locations within each subdivision of the ESA are then processed separately within the geanalysis and access network module.

4.3 Sampling

In order to derive inputs for each geotype used in the CAN module, Analysys has taken a sample of ESAs, stratified by geotype. The size of the sample is currently 200, but this can be changed if needed. This number has been chosen to give reasonable, accurate results on the basis of its calculation time. Currently, the sample can be processed by the geoanalysis and access network module in approximately 4 days. The methodology that was used to design the sample and generate a list of sample ESAs is provided in Annex B.

5 Architecture for the Customer Access Network

5.1 Overview

The CAN is defined as the network that extends from the equipment side of the MDF to the first NTP in the customer's premises. The equipment side of the MDF is included in the Access network, consistent with industry specification.⁷

Analysys has designed a comprehensive set of algorithms that analyse the locations within an ESA and builds a hypothetical fixed network deployment under a set of basic assumptions. As described in section 4.2, the locations are the G-NAF points mapped onto the street network. These are equivalent to the points where the copper cable leaves the street trench to enter the final drop. We call this the "final drop point" (FDP) in the geoanalysis and access network module to distinguish it from the final drop.

The network deployments are encoded in Visual Basic stored within the *Access – Code.xls* workbook, which both takes locations from and stores outputs within the Access DATA workbooks. Two basic network styles have been created: one is an "urban" deployment and the other is a "rural" deployment. Currently, a single type of deployment must be chosen to be used across each geotype.

Urban deployment

Our modern network urban deployment is summarised below in Figure 5.1.

⁷

ACIF C559:2006 unconditioned local loop service (ULLS) network deployment.

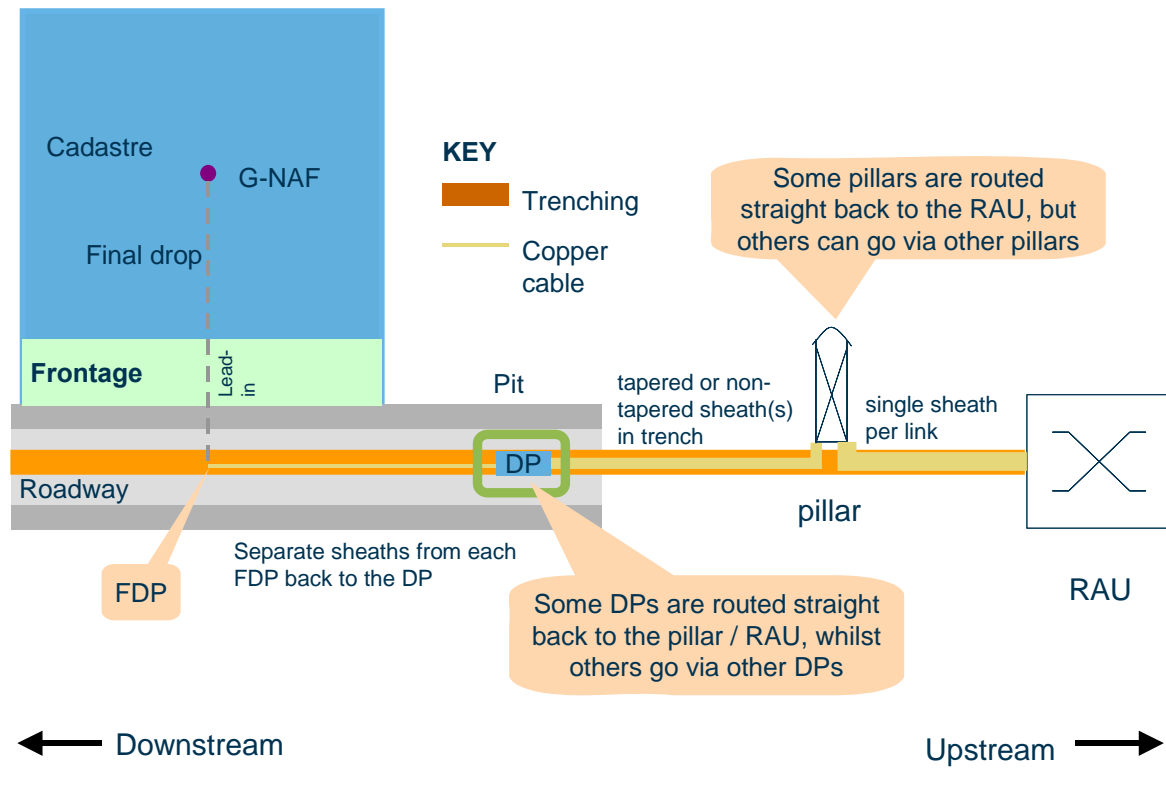


Figure 5.1: Modern urban deployment [Source: Analysys]

Points of demand are grouped into clusters that are served by a distribution point (DP): this is called a “DP cluster.” DPs are grouped into larger clusters that are served by a pillar: these are called “pillar” clusters. Each DP cluster has one point that is designated as the DP serving that cluster. Similarly, each pillar cluster has one DP that is designated as the pillar. The one exception is the pillar cluster which contains the remote access unit (RAU), in which case the actual RAU location is chosen as the central point.

The first stage is to deploy trench to join the points within each DP cluster back to the designated DP. The second stage is to identify a route for each DP back to its parent pillar. This route can possibly be via other DPs and will use both additional trench as well as existing trench within the DP clusters wherever it is seen as more cost-efficient. The cable deployed along these routes between DPs can either be tapered or non-tapered. Finally, each pillar is then routed back to the RAU, possibly via other pillars, using the existing trench deployed within the pillar clusters wherever it is seen as more cost-efficient. Each pillar-RAU link uses an individual cable. Where the link between a pillar and the RAU would lead to copper loop lengths that are too long, an LPGS system is installed at the pillar and linked back to the RAU with fibre.

A fibre deployment can be used to serve locations with particularly high demand. This either uses:

- a fibre ring to connect all pillars back to the RAU
- a fibre ring to connect those pillars with existing locations of high demand back to the RAU
- or direct point-to-point links back to the RAU.

Locations are served by fibre if their demand exceeds a certain level. Such locations are assumed to have a single copper line serving them, which allows a trench network to be derived between the location and its parent pillar. A fibre cable is then laid along this trench to connect the point back to its parent pillar. In the first two cases, the fibre ring is derived separately. In the third case, an individual fibre cable is linked back to the RAU directly from the pillar, via the existing trench for the copper line back to the pillar.

The geoanalysis and access network module is capable of deploying either a tapered or a non-tapered cable network within pillar clusters. This allows comparison between these two options. Both cases dimension the cable by being aware of the demand ‘downstream’ at all DPs in the pillar clusters. The downstream demand at a DP is defined as the demand whose path back to the pillar passes through that DP.

We have run the geoanalysis and access network module assuming both non-tapered and tapered deployments and identified that, for most geotypes, a non-tapered architecture can be deployed at lower cost. The main issue for the tapered architecture is the higher opex for jointing that arises from the increased likelihood of faults occurring. A summary of our comparison is provided in Annex F.

The default assumption used in the geoanalysis and access network module is a non-tapered cable deployment in all geotypes.

Rural deployment

Our modern network rural deployment is summarised below in Figure 5.2.

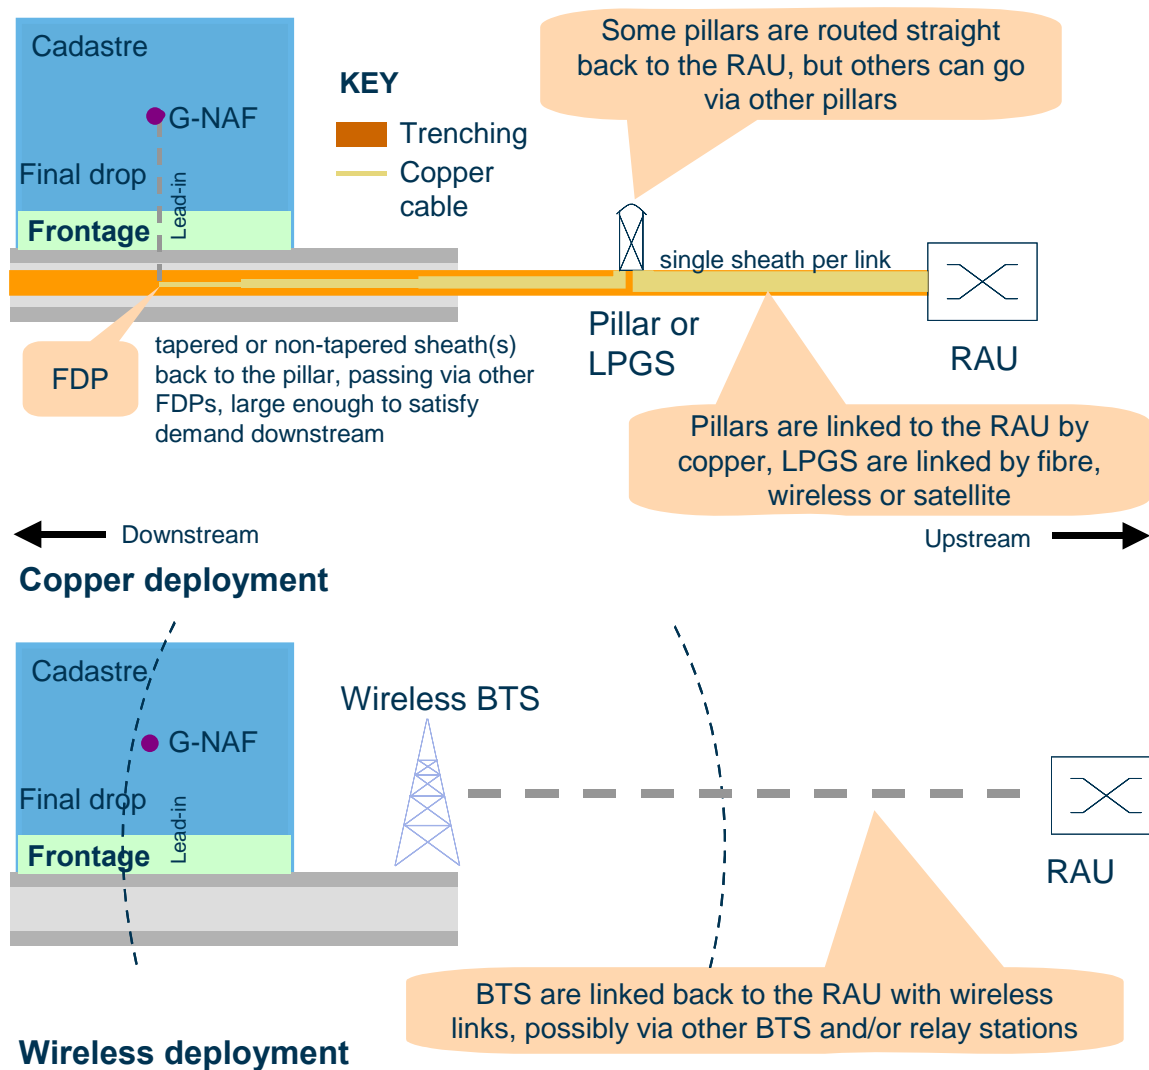


Figure 5.2: Modern rural deployment [Source: Analysys]

Using a cost-based decision, explained in section 5.4.1, points of demand are determined to be serviced by either copper or wireless deployments. Copper-fed points of demand are grouped into clusters that are each served by a pillar⁸. A point within each pillar cluster is designated as the pillar for that cluster, unless it contains the original RAU location. In this case, the RAU is chosen as the central point. When constructing the copper network, the FDP locations are used.

A trench network is designed within each pillar cluster to join the points together. Cable within each pillar cluster can be either tapered or non-tapered. Each pillar is then routed back to the RAU, possibly via other pillars, using an existing trench where possible. Where the link between a pillar and the RAU would be too long for effective service, an LPGS system is installed at the pillar. The backhaul is determined by a cost-based decision, from a choice of either fibre, wireless or satellite.

⁸ Note that the "DP cluster" stage is not used in rural deployments. It is assumed that the buildings in these rural areas are more dispersed and will therefore be served by their own pit.

Relay stations can also be used to lengthen wireless links back. Satellite backhaul is used if the wireless link would require too many relays. If a wireless link is used, then the trench and cable calculated for the link are removed.

The wireless-fed points are also clustered, but into areas that can be served by a base transceiver station (BTS). This clustering uses the original G-NAF coordinates as a better estimate of the actual location of the customer rather than the street point, since the final drop is not relevant for wireless.

It is currently assumed that points of high demand within pillar/large pair gain system (LPGS) areas are connected directly to their parent node and back to the RAU with fibre. It would be unusual to use a fibre ring in a rural deployment, so we use a point-to-point topology. These locations with high-demand are assumed to have a single copper line serving them, which allows a trench network to be derived between the location and its parent pillar. A fibre cable is then laid along this trench to connect the point back to the pillar and then the RAU.

The geoanalysis and access network module is capable of deploying either a tapered or a non-tapered cable network within pillar clusters. This allows comparison between these two options. Both cases dimension the cable by being aware of the demand ‘downstream’ at all points in the pillar clusters. The default assumption used in the geoanalysis and access network module is a non-tapered cable deployment in all geotypes.

The remainder of this chapter sets out the description of these algorithms as follows:

- Section 5.2 describes the basic building block algorithms used in both deployments.
- Section 5.3 describes the traditional deployment that is assumed in urban deployments in more detail.
- Section 5.4 describes the traditional deployment that is assumed in rural geotypes.

5.2 Basic algorithms

Although there are two distinct deployment methodologies, the same fundamental algorithms are used in both: these are clustering algorithms and spanning tree algorithms, as explained below. A simplified version of the TSP algorithm described in section 7.7.2 and used for the core network architecture is used when constructing fibre rings.

5.2.1 Clustering

We use clustering algorithms to group locations into groups to be served by a network node. Our clustering algorithm requires two inputs:

- *Capacity criterion* – if each point has an associated demand, then do all of the points taken in the cluster together exceed that capacity?
- *Distance criterion* – if we calculate the centre of the points in a cluster, or determine a central point in the cluster, then do all points in the cluster lie within a certain distance of that centre?

Given a set of locations and these basic inputs, clustering passes through two stages: a creation phase and a refinement phase. The creation phase determines the number of clusters using a top-down methodology, starting with the initial set of all locations. The refinement phase uses a sequence of individual algorithms to improve these initial clusters according to various criteria. These are summarised in Figure 5.3 below.

<i>Algorithm</i>	<i>Description</i>
Simple reassignment	Moves a point from one cluster to another cluster subject to the constraints, so that the distance from the point to the cluster centre is minimised
Full optimisation	Moves a point from one cluster to another subject to the constraints, so that the total distance within a cluster between its points and cluster centre is minimised
Swap	Swaps two points sitting in different clusters, subject to the constraints, to minimise the distance between these points and their cluster centres
High-demand simple re-assignment	Moves a point with at least two units of demand from one cluster to another cluster subject to the constraints, so that the distance from the point to the cluster centre is minimised
High-demand swap	Swaps points which have at least two units of demand from one cluster to another subject to the constraints, to minimise the distance between these points and their cluster centres
Single demand swap	Swaps two points with unitary demand sitting in different clusters, subject to the constraints, to minimise the distance between these points and their cluster centres

Figure 5.3: *Refinement algorithms used within the clustering algorithm [Source: Analysys]*

All of these clustering algorithms make distance-based decisions using only crow-flies distance.

It is important to note that the distance criterion is *not* intended as a technical limit. For example, a distance criterion of 600m for a pillar cluster in an urban ESA is not intended to mean that such pillars cannot serve locations further than 600m away. Rather, it should be seen as a *control* for the algorithm, which allows the refinements to produce more sensible, tightly knit clusters by reducing the number of points for consideration. In urban areas, most pillar clusters can reach capacity within such a distance anyway, due to the density of locations: the capacity constraint dominates in this case.

More detail on the clustering algorithms is provided in Annex C.1.

5.2.2 Spanning tree algorithms

There are a variety of algorithms available in order to determine how trench and cabling are set out in a CAN for a set of locations. These algorithms determine a set of edges linking the locations forming a tree, allowing any point to reach any other point by following the edges of the tree. These algorithms fall into two main groups, based on their output. These are called *minimum spanning trees* and *minimum Steiner trees*.

Minimum spanning tree

These algorithms calculate the shortest/cheapest amount of duct/cable to link a set of locations (with varying demand) together into a tree. The simplest algorithms of this type effectively derive the shortest duct network to link all premises. However, such an output may not be ideal, since focusing on minimising the duct network may lead to inefficient cable deployments. A simple, albeit extreme, example of this is shown below in Figure 5.4.

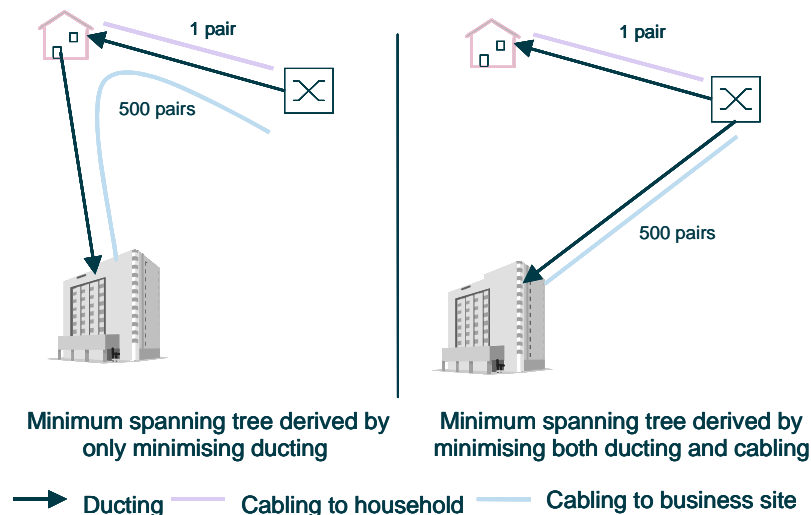


Figure 5.4:
Difference in
optimising duct
alone and both duct
and cabling [Source:
Analysys]

Figure 5.4 shows a very simple case of a household, a large business and a switch. The left side shows the deployment when only ducting is optimised. The shortest route of duct is to connect the business via the house. However, since the business site needs 500 pairs, significantly more cable is required in this scenario than in that on the right, which has a marginally longer distance of duct. Although this is only an illustration, it does demonstrate that, for more complex clusters, the cost of cabling should also be considered.

► *Minimum Steiner trees*

Steiner tree algorithms are an enhancement of minimum spanning trees, in that they permit the creation of new intermediate nodes to be created in order to generate the shortest network architecture. This is analogous to adding an additional access point (e.g. pit) to the existing set of points. The minimum Steiner tree is always shorter than (or in the very unlikely case, equal to) the length of the minimum spanning tree. Calculating minimum Steiner trees for even a sample of

regions of Australia would be highly expensive computationally and therefore unlikely to be feasible for the purposes of this LRIC model.

Although the minimum Steiner tree algorithm is a preferred methodology for deriving the shape of the CAN, its computational expense cannot be ignored. By comparison, there exist tested implementations of minimum spanning tree algorithms that are computationally feasible. We have focused on improving these algorithms and have made several enhancements to improve their effectiveness.

► *Minimum spanning trees implemented*

We have used two minimum spanning tree algorithms in the access network design. Both minimise a cost function which includes the costs of both cabling and trenching when determining the shape of our spanning trees. Although the output may not be a minimum spanning tree, it will be a configuration of lowest cost. Each of these cost functions have been calibrated for our model using a selection of ESAs across the geotypes.

- A modified version of the *Prim* algorithm, similar to that in the FCC's Hybrid Proxy Cost Model (HCPM),⁹ is used to derive the minimum spanning tree for the following elements:
 - trench within DP clusters in the urban deployment
 - trench within pillar clusters in the rural deployment
 - wireless links between base stations, wireless LPGS and RAU in rural deployments.
- A version of *Dijkstra's* algorithm¹⁰ is used to derive a minimum spanning tree to:
 - connect the pillars to the RAU for both deployments
 - connect the DPs to their parent pillars in urban deployments
 - identify the optimal link between any two pillars for the purposes of building a fibre ring.

The proxy cost function for the copper and fibre deployments used in our minimum spanning tree algorithms is of the form:

$$k_1*d + k_2*c + k_3*d*c$$

where d is the distance between two points and c is the number of pairs required between the two points (the capacity). This function estimates:

- *trench* cost in the network with the term k_1*d
- *jointing* costs in the network with the term k_2*c
- *cable* costs in the network with the terms k_3*d*c .

⁹ Bush, C. A., Kennet, D. M., Prisbrey, J. and Sharkey, W. W (October 1999), *Computer Modelling of the Local Telephone Network*, Federal Communications Commission.

¹⁰ E. W. Dijkstra, (1959), *A note on two problems in connexion with graphs*. In *Numerische Mathematik*, 1 (1959), S. 269–271.

The values of k_1 , k_2 and k_3 are inputs in the geoanalysis and access network module and can be changed in *Access – CODE.xls*.

This cost function is employed slightly differently in the two algorithms, but both cases:

- allow the maximum cable distance between RAU and household to be controlled, by considering both the costs of trench and cable
- allow the ability to track the cabling requirements (tapered or non-tapered) within any link, by dimensioning the cable according to the demand that the link needs to be able to serve.

This proxy cost function makes incremental cost decisions when developing the tree, which differs from the global decision that would create the lowest cost spanning tree. Therefore, the cost function has to be calibrated to result in the lowest cost spanning tree. This is why the cost function is a *proxy* cost function rather than an *actual* cost function.

The modified Prim and Dijkstra algorithms are described in more detail in Annex C.2.

5.3 Traditional urban deployments

The urban deployment can be applied to any geotype within the geoanalysis and access network module, but currently is used for geotypes 1-7 and 10. The deployment is constructed as shown below in Figure 5.5 and broadly splits into two phases: clustering and trench/cable network design.

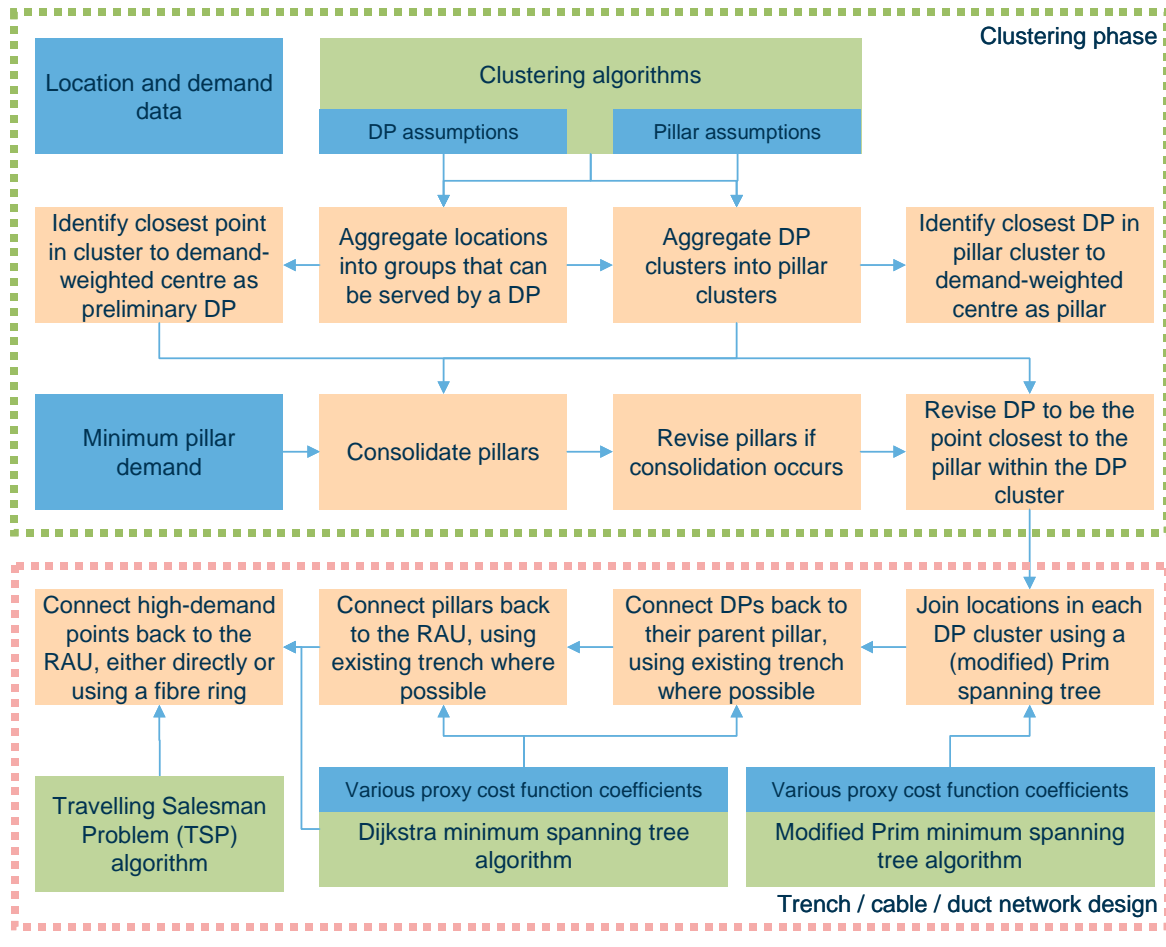


Figure 5.5: Process for determining deployment in "urban" geotypes [Source: Analysys]

The following sections explain in more detail:

- clustering
- trench and cable deployment
- duct deployment
- pit deployment
- jointing.

5.3.1 Clustering

The urban deployment first clusters locations into DP clusters and then groups these clusters further into pillar clusters. More precisely:

- Points are grouped into DP clusters and the point closest to the demand-weighted centre is deemed the DP. The DP cluster containing the RAU uses the RAU location as its DP.

- DP clusters are grouped into pillar clusters using the DP locations. The provisional DP closest to the demand-weighted centre of the cluster is deemed to be the pillar. The pillar cluster containing the RAU uses the RAU location as its pillar.
- DP clusters are grouped into pillar clusters using the DP locations. The provisional DP closest to the demand-weighted centre of the cluster is deemed to be the pillar. The pillar cluster containing the RAU uses the RAU location as its pillar.

Our clustering methodologies consider two types of cluster capacity; absolute capacity and practical capacity (i.e. absolute capacity subject to an assumption of normal utilisation). For example, we assume that pillars have a maximum capacity of 400 SIOs: this constraint is never relaxed. However, we also assume that a pillar cluster is designed to be 90% utilised (i.e. 10% spare), meaning that a pillar cluster should usually not exceed 360 SIOs. However, we do allow this constraint to be relaxed in a process of pillar cluster consolidation.

Having derived the pillar clusters, this consolidation process determines whether any pillars can be merged whilst still obeying the absolute pillar capacity. In the case of the urban deployment, this process will attempt to avoid deploying pillars that serve only a small amount of demand, below a certain critical level. This is done by allowing such clusters to merge with their nearest neighbour and will even relax the distance criterion for the merged cluster to allow this to happen. An example is shown below in Figure 5.6 for Waymouth, which has three isolated clusters with less than five units of demand each that would otherwise each be served by their own pillars.




Figure 5.6:
Isolated points in
Waymouth [Source:
Analysys]

This process still obeys the absolute capacity constraint of the pillar, although it may use some of the assumed spare capacity in order to accommodate the additional demand.

As described in section 5.2.1, refinement processes are used to improve the preliminary clusters. The refinements for each stage in the urban deployment where clustering is used are shown below in Figure 5.7:

Stage	Clustering into DP areas	Clustering into pillar areas	Clustering pillars into fibre rings
1	Simple re-assignment	Simple re-assignment	Simple re-assignment
2	Swap	Swap	Swap
3	Full optimisation	Full optimisation	Full optimisation
4	High-demand simple re-assignment	Simple re-assignment	Simple re-assignment
5	High-demand swap	Swap	Swap
6	Single-demand swap		
7	Simple re-assignment		
8	Swap		

Figure 5.7: Order in which the refinement algorithms are used for each application of clustering in the urban deployment [Source: Analysys]

The high-demand and single-demand refinements have been designed specifically for DP clustering, where the other refinements struggle to improve the clustering due to the small capacities of the DP clusters. These refinements are particularly useful for ESAs in Bands 1 and 2, where the proportion of points with multiple units of demand can be high.

In DP clustering, the maximum absolute capacity is used to limit the total number of SIOs from a collection of locations. Where a single location has a demand higher than that of the maximum absolute capacity, but less than the limit for it to be served by fibre, it becomes a DP cluster of its own with no FDP locations included in that cluster.

5.3.2 Trench and cable

Having clustered the locations into clusters served by DPs and then pillars, the trench and cable networks are built up incrementally, by using the minimum spanning tree algorithms described in section 5.2, joining up the following elements:

- points of demand within DP clusters (FDP-DP links)
- DPs to their parent pillar (DP-DP links)
- pillars to their parent RAU (pillar-pillar links)
- pillars into a fibre ring and then fibre-fed points back to the RAU (either fibre links via the ring or point-to-point links back to the RAU).

FDP-DP

FDP-DP links use the modified Prim algorithm to form the trench network within a DP cluster and calculate the number of copper pairs required to reach the end-customer location. Being at most four or so points, these are calculated very quickly. The (non-repeated) trench required to connect all the points to the DP is calculated. When calculating the copper sheath required, it is assumed that individual cables go from the DP to the FDP: joints are calculated at the FDP dependent on the number of pairs deployed to the FDP. Moreover, only particular cable sizes are used, with 2-pair being the smallest (followed by 10-pair).

It has also been defined how each level of demand at a point is served by combinations of cable. For example, three units of demand are currently assumed to be fed by two 2-pair cables, whilst 6 units of demand are fed by one 10-pair cable.

The length of all cable links in the FDP–DP tree are calculated using Euclidean distance, since at this scale, DP clusters are likely to lie along roads and the distances will almost all be crow-flies. In some cases, the points in a DP cluster may actually lie around a corner and using Euclidean distance will not capture the actual deployment distance of the trench. As a result, we apply the trench-sharing coefficient to the total FDP-DP trench in the ESA as an adjustment, which we describe in section 4.2.5. However, the small size of DP clusters means that FDP–DP trench is not usually a large fraction of the total trench required.

With the assumption of a DP capacity of four units of demand, in almost all cases the cable requirements will only be 2-pair cable between FDPs and DPs. This is because locations with four units of demand or less are only served by one or two 2-pair cables. Although locations with five

or more units of demand would need 10-pair cable, they will always be their own DP, in which case no FDP–DP cable is required.

DP-pillar

The modified Dijkstra algorithm uses a pre-calculated full mesh of trench linkages to join DPs to their parent pillar. The pre-calculation and subsequent algorithm proceed by:

- for each pair of DP clusters, identifying a point within each cluster that produce the shortest link between the two DPs
- creating a fully meshed set of lowest proxy cost linkages between DPs using these points. These links can form a ‘daisy-chain’ via other DPs and assume that the trench within the DP clusters has already been laid
- using the Dijkstra algorithm to select the least proxy cost subset of linkages that allow each DP to be joined back to the pillar.

This process is illustrated for a simple example of three DP areas, one of which contains the parent pillar, in Figure 5.8 below:




Figure 5.8: Use of the Dijkstra algorithm [Source: Analysys]

These linkages can be assumed to have two designs, either tapered or non-tapered. We understand that the existing network implements a non-tapered distribution network and we have implemented what we believe to be a fair approximation of this. Our tapered and non-tapered methodologies are explained below:

Tapered

For a given DP-DP link, a single cable sheath is used, whose size is dimensioned by the demand downstream of the DP furthest along the trench network away from the pillar. For example

- to link an unconnected DP to the network, the cable used is the smallest that can accommodate the capacity of this DP.
- to link a DP to the pillar, the cable used is the smallest that can accommodate the capacity of all the DPs linked to the pillar along this network path.

Jointing occurs at each DP. We also assume a certain utilisation of the distribution network cabling. For example, a utilisation of 100% means that a 30-pair cable can be used for up to 30 units of downstream demand, whereas with a utilisation of 80%, it can only be used for up to 24 units of demand.

Non-tapered

In the non-tapered case, only two cable sizes are used. The main cable size is 100 pairs, with a minor cable size to use at the most extreme points in the network. We also assume a certain level of cable utilisation in this case. Setting the minor cable size to zero means that the distribution network can be fully non-tapered.

The cable size is dimensioned using the downstream demand at the DP as in the non-tapered case. If the downstream demand exceeds the minor cable size (multiplied by the utilisation), then the required number of sheaths of the major cable are used. As an example, if we assume a utilisation of 80%, then 90 pairs are served by 2×100 -pair cables.

The modified Dijkstra algorithm leads to daisy-chaining of DPs, since a DP may go back to the pillar via several DP clusters and will go through each DP. This is assumed to occur because any jointing would require access to the cable and a DP is the lowest cost access point for this to occur.

Since the cable must pass through the DP location of every DP cluster that it passes through back to the pillar, this can lead to a small amount of “back-tracking” in the route, as demonstrated below in Figure 5.9. This will lead to a small over-estimation of the cable requirements. However, since this is restricted to the scale within a DP cluster, we do not believe this to be a significant source of inefficiency in the route.




Figure 5.9: Demonstration of doubling back by the DP-pillar cable [Source: Analysys]

Pillar-RAU

► Cable routing

A similar algorithm is used as in the DP-pillar case, but with some differences:

- The proxy cost function used has different cost coefficients, which is necessary as the main cable may be specified with more pairs than the DP-pillar cable, which allows a different bias between cable and trench in the pillar-RAU links.
- The link between a pillar and the RAU is assumed to be a discrete cable: if two pillars use the same trench to get back to the RAU, then it contains separate sheaths.
- A separate duct is used along the pillar-RAU link routes and is assumed to have a fixed capacity of pillar-RAU cables per duct.

Due to the complexity of calculations at this level, there are several simplifying assumptions that have been made compared with the DP-pillar level. These are:

- Pillar linkages are calculated by finding the two closest *DPs* in the pillar clusters, not the two closest *points* (and the link is assumed to be a single trench linking the two *DPs*, rather than passing through any points in their clusters). This may over-estimate the trench required. However, a point in the *DP* cluster cannot be more than 50m from the cluster centre, so the net error will be of the order of several metres of trench per pillar linkage in an ESA.
- If a pillar needs to pass through several other pillar areas to get to the RAU, then it is assumed to pass through each pillar location. This may over-estimate the cabling required, since a shorter route is likely to be possible without going through each pillar location.

These assumptions simplify the algorithm significantly at this stage. The algorithms also remove any occurrences where the cable doubles back on itself on its way to the RAU, which can be a more significant effect at the scale of the pillar cluster. We re-calculate the required ducting, sheath lengths and jointing respectively.

Having derived the cable length to get back to the RAU, if this leads to locations having loop lengths that are too long to be linked by copper, then an LPGA is assumed to be installed. The pillar-RAU copper sheath/pair length is replaced with a corresponding fibre sheath/cable length. We assume that LPGA in the urban deployment are never linked to the RAU using wireless or satellite backhaul. Copper demand served by LPGA are unavailable for ULLS.

Based on the current maximum loop-length threshold, our deployed network currently leads to about 5% of demand in Band 2 being served by LPGA, which is very similar to actual levels.

Our assumed maximum permissible loop length take into account the performance levels of the cable gauges that are available.

► *Cable gauge*

The performance of copper loops for providing fixed services can vary significantly depending on both the length of the loop and of the gauge (thickness) of the copper used. Shorter loops can supply services using a thinner gauge than longer loops. In our model, we have assumed that three different gauges can be deployed in order to provide ISDN/ADSL services, as shown below.

<i>Gauge (mm)</i>	<i>Loss per km for ADSL (dB)</i>	<i>Loss per km for ISDN (dB)</i>
0.40	13.81	8.25
0.64	8.62	4.03
0.90	6.36	2.48

Figure 5.10: Types of cable deployed

[Source: ACIF

C559:2006

Unconditioned Local

Loop Service (ULLS)

Network Deployment]

We have used the outputs of the geoanalysis and access network module to calculate the distribution of gauges that can be used in order to satisfy the constraints of ADSL (maximum loss of 59.39dB, tested at 300KHz) and/or ISDN (maximum loss of 45.45dB, tested at 40KHz).

Distributions of gauge deployment are calculated for the copper within clusters served by pillars, within clusters served by LPGA and from pillars back to the RAU separately. These distributions can be found in the Cost module and are used to derive blended unit costs for the cable sheaths deployed within the access network.

Fibre ring

In Bands 1 and 2 in particular, there are locations with very high demand, which would require significant copper cable in the DP cluster. In reality, such points would be fed by fibre. For each geotype, the geoanalysis and access network module assumes that fibre locations are fed by either:

- a discrete fibre cable from the location to the parent pillar where it is optically joined to a fibre ring going through every pillar and the RAU

- a discrete fibre cable from the location going back to the RAU via its parent DP and pillar.

We understand that the existing network uses fibre rings in Band 1 and point-to-point otherwise. Therefore, we assume that geotypes 1 and 2, being areas of the most concentrated demand, use fibre rings to increase the chance of having assets near when that demand arose. The rings include the RAU and use nodes at each pillar location. We assume that the remaining geotypes use point-to-point fibre links.

An FDP is assumed to be fed by fibre if it has more than a certain level of demand. Regardless of this, a single unit of demand is assumed to be fed by copper. This is so a path back to the DP and pillar can be derived using the minimum spanning tree algorithm. A fibre cable is laid along the entirety of this path from the FDP back to the pillar.

In order to link these locations back to the RAU using point-to-point links, the fibre cable is assumed to extend back, following the pillar-RAU link used in the copper deployment.

When using fibre rings, the pillar locations are clustered so that each pillar lies in a cluster and each cluster contains the RAU. One ring is constructed per cluster using a simplified version of the TSP algorithm used in the core network deployment. The fibre cable requirements of the ring itself are dimensioned by the number of fibre-fed locations lying on the ring.

5.3.3 Duct

Duct calculations are undertaken for all ESAs. For each geotype, these calculated results are used to extrapolate the duct requirements of the whole geotype for the purposes of the CAN module.

Six types of duct are identified within the network, depending on the purpose of the cable that they are carrying. The types of duct and the number of cables that they are assumed to carry are summarised below in Figure 5.11:

<i>Nature of link that duct is used for</i>	<i>Dimensioning of duct</i>
FDP–DP	At most 1 duct per link
DP–pillar (intra-pillar)	4 cable sheaths per duct
Pillar–RAU (inter-pillar)	2 cable sheaths per duct
LPGS–RAU	100 cable sheaths per duct ¹¹
Fibre-fed FDP–pillar (intra-pillar)	100 cable sheaths per duct ¹²
Fibre rings, or for point-to-point fibre back from pillar to RAU (inter-pillar)	100 cable sheaths per duct ¹²

Figure 5.11: Types of duct deployed

[Source: Analysys]

¹¹ These capacities will effectively dimension at most 3 ducts in a single trench for fibre: at most 1 for LPGS–RAU, at most 1 for point-to-point fibre from FDPs back to the pillar and at most 1 for point-to-point fibre from pillars back to the RAU. These three different uses of fibre will not share the same duct.

Duct is laid along the full route described for the link. For example, for a cable link between a pillar and the RAU, the duct is laid along the incremental pillar–RAU trench, the intermediate DP–DP trench links in the links between pillars and the FDP–DP trench used in the links between DPs.

We only include a dedicated FDP–DP duct in those FDP–DP links where no other duct is present. Otherwise, the cabling between an FDP and its parent DP uses the local serving pit to property boundary architecture as described below in section 6.3.

The total number of duct required is written, but the number provisioned is taken to be the first number that is larger than that required from the list 1, 2, 4, 6, 8, 12, 16, 20, 24 and 28. The maximum number of duct provisioned is taken to be 28.

This allows, for example, the separation of the LPGS–RAU cable and duct, since this is considered to be part of the core network.

5.3.4 Pits

Pit calculations are completed for all ESAs. As in the case of duct, the calculated results are used to extrapolate the duct requirements of the whole geotype in the CAN module. Analysys has deployed pits in the modelled network similar to those currently in use. Pits are firstly deployed:

- at cluster nodes (DPs and pillars)
- along trench links in the spanning tree that are long enough to required additional jointing between the endpoints.

Six types of pit are used, as shown below. The type of pit deployed at the cluster nodes is the smallest that can accommodate both the number of ducts entering the node and the number of links entering/leaving the node. The dimensioning rules are summarised below in Figure 5.12:

<i>Pits</i>	<i>Pit capacity in ducts</i>	<i>Pit capacity in links</i>
P5	1	1–2
P6	2–3	3
P9	4–5	4 and above
PF12	6–15	4 and above
PF20	16–23	4 and above
PF28	24 and above	4 and above

*Figure 5.12:
Dimensioning of pits
for the cluster nodes
[Source: Analysys]*

In addition, every pillar must have at least a P9 pit present. So, for example, a pillar point which is the intersection of 3 links, with 22 ducts entering it in total would require at least:

- a PF20 pit by the number of duct
- a P6 pit by the number of links
- a P9 pit since it is a pillar.

Therefore, in this case, a PF20 pit is deployed.

Within the geoanalysis and access network module, additional pits are deployed to accommodate additional jointing required wherever trench links in the DP–pillar network exceeds 100m in length and incremental pillar–pillar trench exceeds 250m¹² in length. Therefore, a trench in a DP–pillar route that is 240m long would have two extra pits, whilst the incremental trench that is 600m long in a pillar-RAU link would have two extra pits. These two types of extra pit are included in the cost calculation, but we have not identified their positions.

Within the CAN module, additional ‘serving pits’ are added within DP clusters in order to account for both ‘isolated’ FDPs (i.e. FDPs which the 100-pair DP–pillar sheath does not pass through) and for those DP clusters identified as having a road crossing.

5.3.5 Jointing

We assume that no jointing occurs in the link from an FDP back to its parent DP.

In the copper feeder network between the pillar and the RAU, a full 400 pair joint is added every 500m, based on the current input parameter. This is assumed to be done at the nearest DP within the pillar’s trench network.

No jointing is calculated for fibre cabling deployed, as the cost of fibre access is not as critical to the declared access services. Fibre jointing costs should be included with the fibre material and installation costs in the cost inputs.

Jointing in the distribution network between DPs and pillars is dimensioned in different ways depending on whether tapering is used or not.

Tapered

In the tapered case, jointing at a DP is the sum of the capacity of the demand at the DP and the number of pairs in the downstream cabling. Where a DP has more than one incoming branch, a branching kit for each extra branch is also deployed. This is shown below in Figure 5.13 for a series of DPs lying on a daisy chain back to the pillar, assuming 100% utilisation of the cable.

¹² A value of 250m is used since pillar–RAU cable is deployed in a mix of 0.4mm, 0.64 mm and 0.9mm gauges. We understand that 0.9mm can be hauled for at most 250m, meaning that there should be an access point at least this frequently.




Figure 5.13: Tapered jointing between DPs and the pillar [Source: Analysys]

Non-tapered

In the non-tapered case, jointing at a DP is taken to be just the demand at that DP. Where a DP has more than one incoming branch, a branching kit for each extra branch is also deployed. Finally, additional joints are required in each link wherever the cable distance from the pillar exceeds a multiple of 500m. This is assumed to occur at the DP in the link closest to the pillar.

This is shown below in Figure 5.14 for a series of DPs lying on a daisy chain back to the pillar, assuming 100% utilisation of the cable.

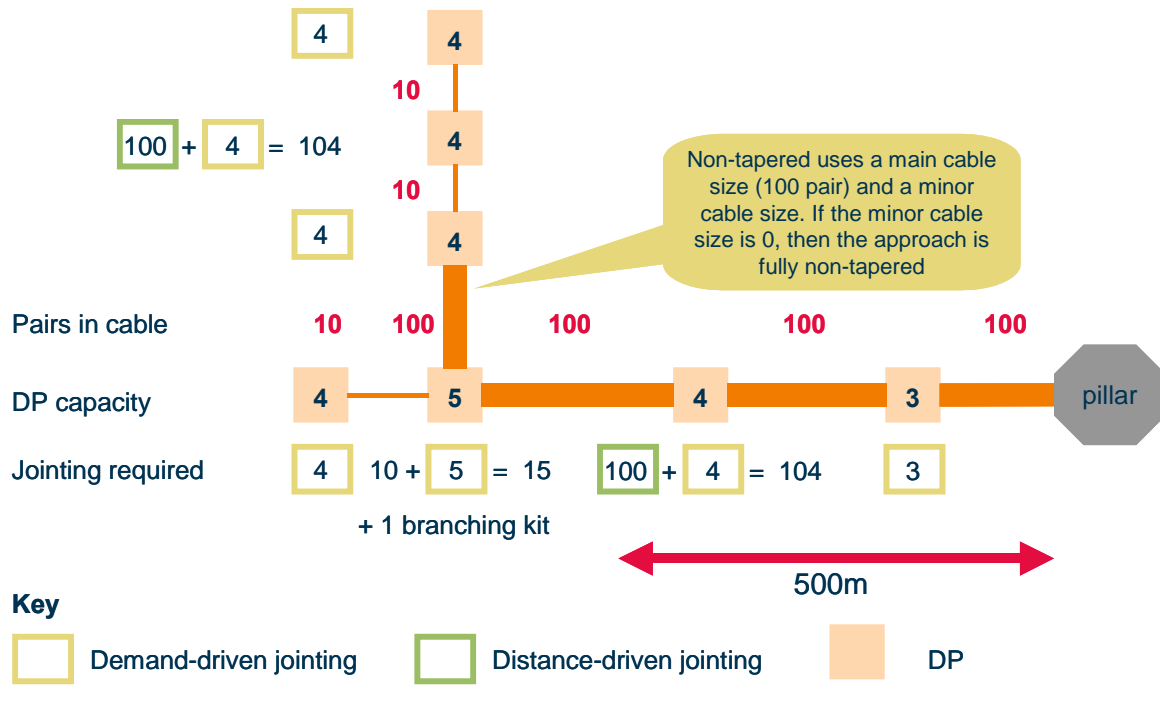


Figure 5.14: Non-tapered jointing between DPs and the pillar [Source: Analysys]

5.4 Traditional rural deployments

The rural deployment can be applied to any geotype within the geoanalysis and access network module, but is currently used for geotypes 8-9 and 11-14. The deployment is constructed as shown below in Figure 5.15, but broadly splits into four phases: allocation of points to copper or wireless deployment, clustering, copper trench/cable/duct network design and wireless network design. The following sections explain in more detail:

- if locations are served by copper or wireless
- clustering
- trench and cable deployment
- duct deployment
- pit deployment
- jointing
- wireless backhaul.

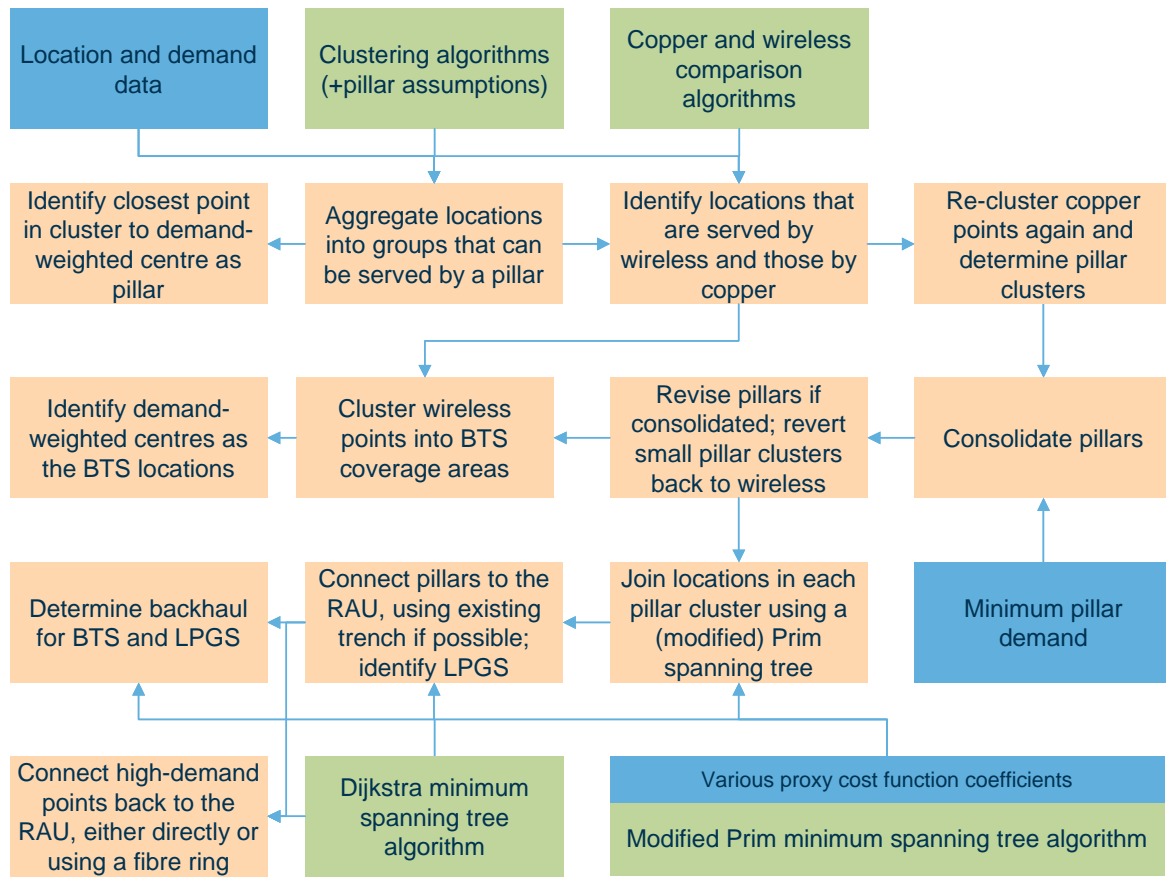


Figure 5.15: Process for determining deployment in “rural” geotypes [Source: Analysys]

Figure 5.15 above details the process involved in the rural algorithm. All of the locations are fed through an algorithm that makes a cost-based decision whether to serve each of them with copper or wireless solutions.

The two sets of locations are then processed separately:

- A copper deployment similar to that in the urban deployment is executed on those locations determined to be served by copper. Either pillars or LPGS can be deployed, but LPGS can use either fibre or wireless backhaul.
- The remaining points are clustered, using different constraints, to be served by a set of wireless base stations.

A backhaul network back to the RAU is then defined for the nodes which require wireless backhaul.

5.4.1 Identifying a location as being served by copper or wireless

ESAs in rural areas can have locations that are too remote to serve with copper at an efficient cost. For the locations in any ESA, we have developed a cost-based methodology to determine those which should be served with copper or fibre and those which should be served by wireless. Cost approximation functions are used to estimate the incremental cost of serving a location with copper or wireless.

The output of this copper-wireless decision is to designate each location as being served either by copper or wireless. The locations stated to be served by wireless are then processed by a rural copper deployment.

There is a central process that is applied to the locations. This algorithm is summarised below in Figure 5.16.

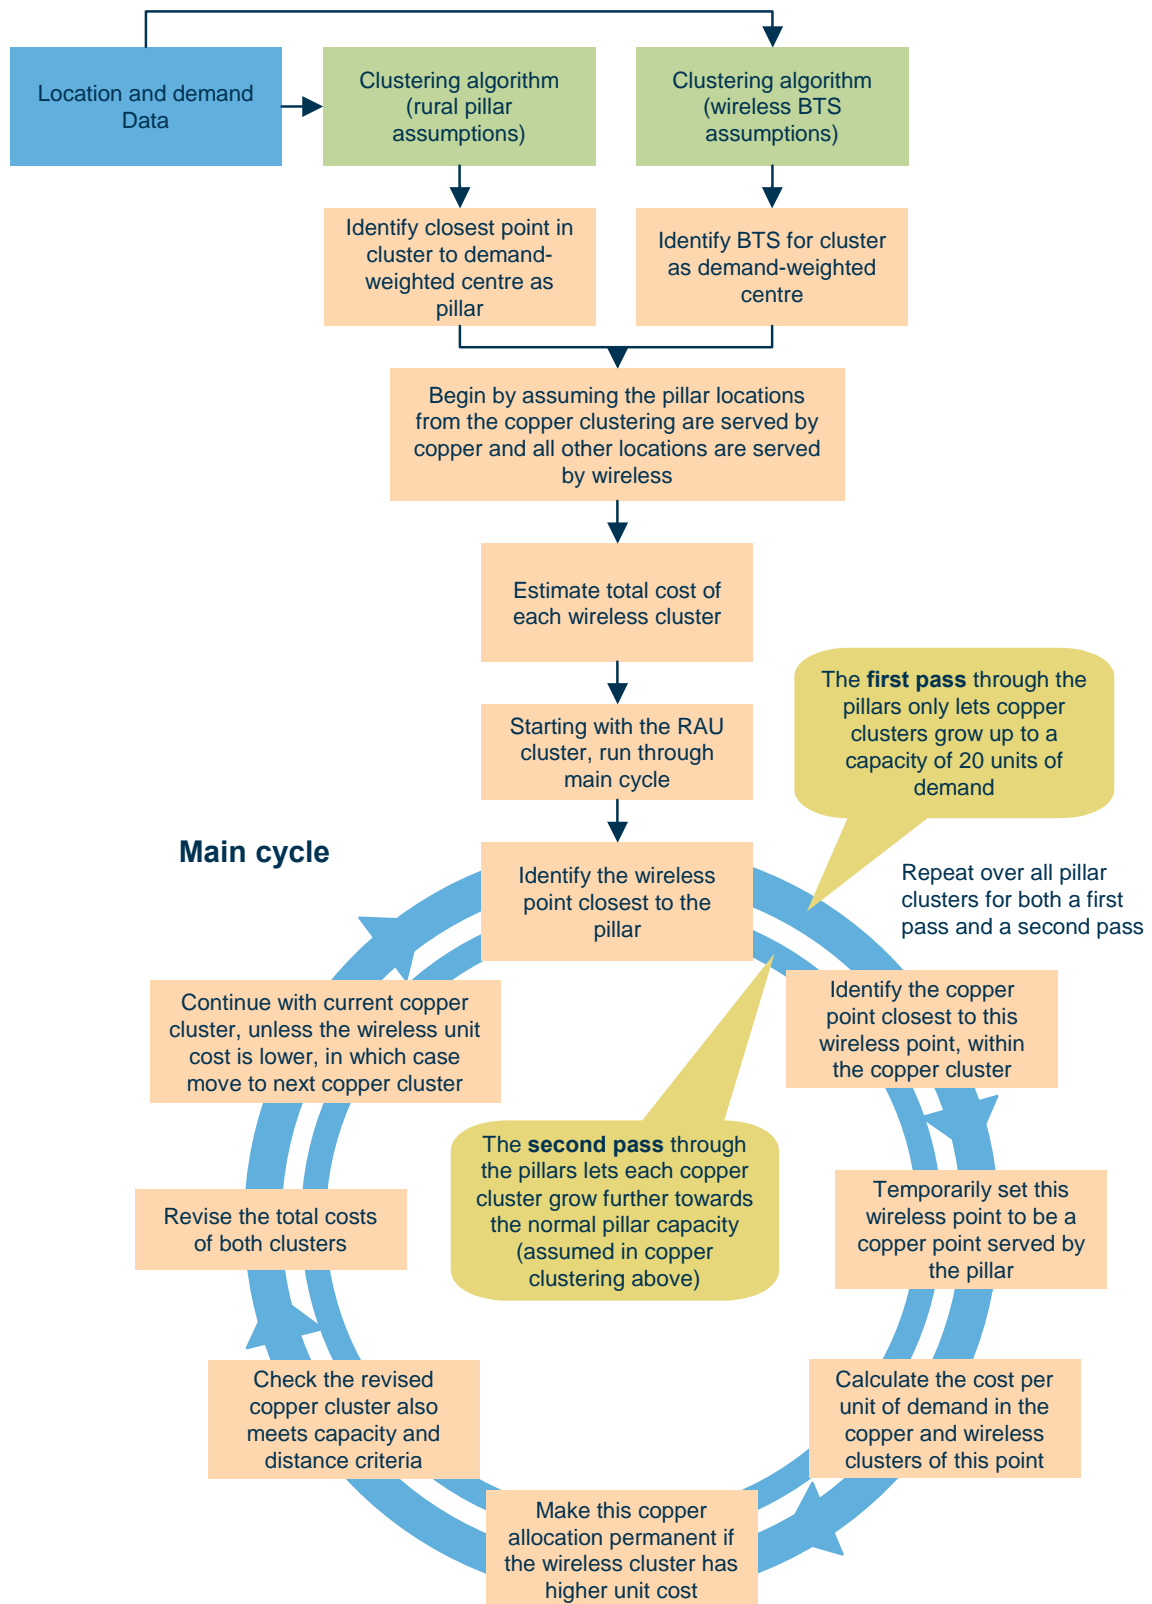


Figure 5.16: Central process for the copper-wireless decision [Source: Analysys]

The cost approximation of a wireless cluster is calculated by:

Base transceiver station (BTS) setup cost + (number of wireless locations served × incremental cost per subscriber)

The cost approximation of a copper cluster is calculated cumulatively during the algorithm using:

New copper cluster cost = Old copper cluster cost + (distance of closest point in the pillar cluster to the unattached point × incremental cost per m)

The coefficients currently used for these two cost functions are based on costs provided by the ACCC, a universal service obligation (USO) report commissioned by the ACMA (previously ACA)¹³ and Analysys industry benchmarks. These are summarised below in Figure 5.17.

<i>Cost category</i>	<i>Relative set-up cost (AUD)</i>	<i>Relative incremental cost (AUD)</i>
Copper (via pillar)	2000	11 per metre
Copper (via LPGS)	80 000	11 per metre
Wireless	200 000	825 per location

Figure 5.17: Coefficients used in the copper and wireless cost functions [Source: ACCC, Analysys assumptions]

Each pillar is considered in turn during the algorithm. It considers trying to transfer wireless-fed points to the copper cluster of the pillar until all remaining candidate points can be served more cheaply by wireless rather than copper. Only then does it move on to try and grow the cluster associated with the next pillar. Moreover, the first pillar considered is the RAU location, which is normally surrounded by the highest density of locations. As a result, if the algorithm was left to add more points to the copper clusters without any control, then it would be likely that it would just develop a single large copper cluster around the RAU and nowhere else.

In order to give additional pillars the opportunity to grow their clusters in addition to the RAU, the pillars are considered in turn twice. The first time that a pillar is considered: it only grows pillar clusters up to a small minimum ('critical') capacity of 20 units of demand. This is an assumption based on the smallest level of demand that a pillar would be deployed to serve. The second time that a pillar is considered, the clusters are permitted to grow up to the capacity limit used in the normal clustering algorithm.

For the remainder of this description, we refer to this process of checking through each pillar in turn twice as the *main cycle*.

Due to the sparse nature of the locations in rural ESAs, LPGS are likely to be deployed in these ESAs. We therefore consider two different *copper deployment scenarios* in parallel, running our main cycle for both of them. These two deployment scenarios differ in terms of the distance constraint used in the calculations for the copper clusters around the pillars.

¹³ *Forward Looking Technologies for the USO 2000–2003*, Gibson-Quai PTY LTD, August 2000.

The *standard deployment scenario* sets the distance constraint for the pillar as a function of the maximum copper loop length back to the RAU:

- If the p–function distance from the pillar to the RAU is greater than this limit, then it is assumed that an LPGS is used and the cluster distance criterion is set to be equal to the maximum loop length.
- If the p–function distance from the pillar to the RAU is less than this limit, then it is assumed that the cluster distance criterion is equal to the maximum loop length less the p–function distance from the pillar to the RAU.

The *‘extended reach’ deployment scenario* assumes that all pillar locations use an LPGS and hence are able to test for copper locations using the maximum copper loop length criteria.

The lowest cost solution of these two iterations is used by the remaining stages of the copper-wireless decision, which clean up the solutions, so that there are no pillar clusters:

- lying within a larger pillar cluster; otherwise the two clusters are merged
- which are smaller than the minimal critical capacity; otherwise they are served by wireless
- that are not convex i.e. there are no copper locations that are further from their parent pillar than a wireless location.

In addition, the RAU always survives this decision as a copper-fed location, unless there are no copper-fed locations at all. This is to be consistent with the scorched-node assumption that we use in our approach.

In summary, the copper-wireless decision goes through two iterations and a series of refinements in order to arrive at a mapping of locations between copper and wireless. This is shown below in Figure 5.18:

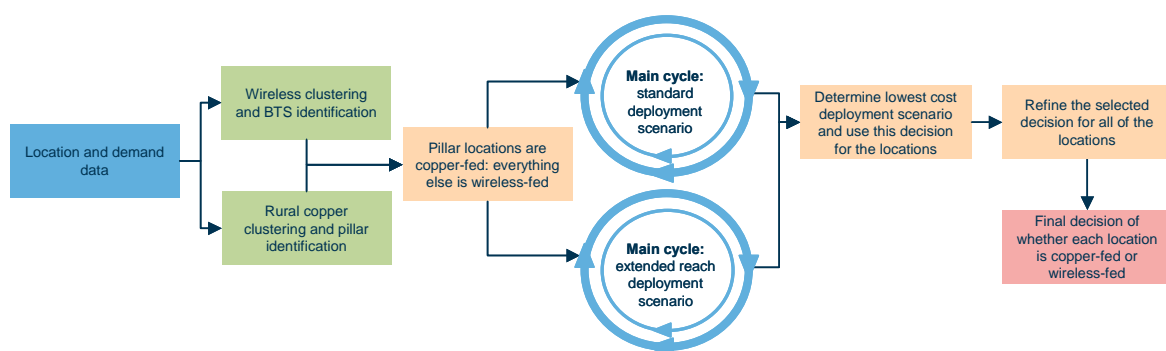


Figure 5.18: Structure of the whole copper-wireless decision [Source: Analysys]

5.4.2 Clustering

The copper-wireless decision described in the previous section designates each location as being served either by copper or wireless. For the purposes of the actual copper deployment, the points that are designated to be served by copper are re-clustered before the trench and copper network are designed. Wireless locations are clustered separately according to the wireless constraints. A wireless backhaul network is derived, if required, in order to link nodes with wireless backhaul back to the RAU.

Copper clustering

The significant difference between urban and rural copper deployments is that rural deployments do not use DP clusters. Therefore, copper-fed locations are clustered directly into pillar areas. The pillar location is assumed to be the demand-weighted centre of all the points in the pillar cluster. The pillar of the cluster containing the RAU is assumed to lie at the RAU.

Since the clusters involved in the rural deployment are much larger, the high-demand and single-demand refinement algorithms are not used.

Wireless clustering

We have assumed that wireless locations are identified by their G-NAF point rather than their FDP on the road network, since the wireless solutions do not have a final drop like a copper solution. The distance criterion is assumed to be the coverage radius of a wireless BTS. The inputs are currently calibrated to 900MHz GSM technology, due to it being a mature, low-cost wireless solution with large coverage radii.

The BTS is assumed to be located at the actual demand-weighted centre of the points in the cluster, not the nearest point in the cluster to that centre.

Satellite clustering

There will be certain BTS that are serving only a handful of locations. As a result, it will be more cost effective to serve these locations by satellite. For each wireless cluster, the algorithm compares the cost of serving the locations by:

- wireless, calculated by the cost of the tower and the cost of one customer premise equipment (CPE) per wireless location
- satellite, calculated by the cost of serving a location by satellite multiplied by the number of locations in the wireless cluster.

Those clusters that can be more cheaply served by satellite have their BTS removed and are set to just be served by satellite. We currently assume a relative satellite cost of AUD8300 per location,

based on costs provided by the ACCC and a USO report commissioned by the ACMA.¹⁴ The wireless costs used are the same as those employed in section 5.4.1.

5.4.3 Copper trench, cable and duct network

The subset of copper-fed locations are then treated in a similar way to the urban deployment, with a spanning tree designed to join each cluster of points back to their parent pillar.

FDP-Pillar

Rural deployments do not use DP clusters, but assume that each location has its own DP. Copper-fed locations are clustered directly into pillar areas and the pillar location is assumed to be the demand-weighted centre of all the points in the pillar cluster, unless the pillar cluster contains the RAU location. The points in a pillar cluster are joined by a minimum spanning tree generated using the modified Prim algorithm, as described in section 5.2.2. Given the much larger clusters being used, this process is computationally more costly and requires more time per location than in its application in the urban deployment.

This tree will be assumed to use cable on either a tapered or a non-tapered basis. Both include an assumption of the level of utilisation of the cable and dimension the cable requirement in a link in terms of downstream demand. The non-tapered case uses a main cable size (100 pair) and a minor cable size. As soon as the downstream demand exceeds the minor cable size (considering the utilisation percentage), then a number of sheaths of the major cable size required are deployed (which also considers the utilisation percentage). For example, assuming a cable utilisation of 75%, a link with a downstream demand of 180 requires three 100 pair cables.

As the tree is grown, the cabling requirements in each link are updated each time. When an unattached point is connected, for each link in its route back to the pillar, the cabling in the link is checked and updated if necessary to accommodate the extra demand from the unattached point.

Pillar-RAU

► *Cable routing*

Like the urban algorithm, the pillar-pillar linkages are determined using Dijkstra algorithm to select from the full mesh of possible solutions. In the rural case, possible linkages are assessed between any two points in separate pillar clusters.

¹⁴ Gibson Quai PTY LTD (August 2000), *Forward Looking Technologies for the USO 2000 – 2003*, http://www.acma.gov.au/webwr/telcomm/universal_service_regime/usotechnology_report2000.pdf.

Having identified the lowest cost cable distance back to the RAU for a pillar, it is then checked to see if it means that the total copper loop length exceeds the maximum limit. If it does, then an LPGS is installed in its place. Unlike the urban deployment, LPGS in rural deployments can be served by either fibre, wireless or satellite. Choosing between a fibre or wireless link is determined by a cost-based decision.

The proxy cost of the fibre link is given by

$$k_1 * D_T + k_2 * c * D_C$$

where D_T is the incremental trench required between the LPGS and the RAU, c is the number of fibres in the cable required between the two points, D_C is the fibre sheath length required and n is the number of regenerators required. It is assumed that very long fibre links require a regenerator at distance intervals of 90km: in practice these links are unlikely to be deployed within an ESA.

The proxy cost of the wireless link is given by

$$k_1 + k_2 * M + k_3 * n$$

where M is the cost multiplier of backhaul capacity required between the two points and n is the number of relay stations required. It is assumed that wireless links require a relay station at regular intervals in order to boost the signal. We currently assume that three backhaul capacities are available: 2Mbit/s, 8Mbit/s and 34Mbit/s. Backhaul capacity is aggregated on the units of copper demand of the locations within the wireless cluster. A point with a single unit of demand receives a residential backhaul capacity, whilst a point with multiple units of copper demand receives a business backhaul capacity per unit of demand. We have made the simplifying assumption that the cost of the relay station does not vary with the backhaul capacity, since the larger backhaul capacities are seldom used in practice.

Of these two options, the one with the least cost is then used to backhaul the LPGS. However, if a wireless link is used which would require more than a specified upper limit of relay stations, then satellite is used instead.

As in the urban case, we have also applied post-processing to remove any occurrences where cable in this part of the network doubles back on itself back to the RAU. We re-calculate the required ducting, sheath lengths and jointing respectively.

► *Cable gauge*

We use the same methodology to calculate the distribution of cable gauges deployed in the rural deployment as we use in the urban deployment.

Fibre deployments

High-demand points are very sparse in the more rural ESAs in Bands 3 and 4. We have assumed that point-to-point fibre is used as described in the urban deployment in section 0.

5.4.4 Duct

All types of ducts deployed in the urban case are also deployed in the rural case, except FDP–DP duct, which does not exist. The dimensioning rules used in their deployment are the same as those in the urban deployment.

5.4.5 Pits

Pits are deployed at each FDP in the copper clusters, since each FDP is taken to be its own DP in the rural case. The dimensioning rules used in their deployment are the same as those in the urban deployment, with the exception that no additional pits are accounted for in the CAN module. This is because the criteria for additional pits do not apply in rural deployments, since every location already has its own provided.

5.4.6 Jointing

Like the urban algorithm, the jointing required in the links within the copper network are dimensioned using both a demand-driven component and a distance-driven component. The number of branching kits needed at each junction in the copper network are also determined.

5.4.7 Wireless backhaul

For rural deployments that deploy wireless BTS and/or LPGS with wireless backhaul, a wireless backhaul network is created back to the RAU. This is created using the modified Prim algorithm on these nodes. The proxy cost function used is of the form:

$$k_1*d + k_2*M + k_3*n$$

where d is the crow-flies distance between the two nodes, M is the cost multiplier of backhaul capacity required between the two points and n is the number of relay stations required. This is almost the same form as that for the wireless backhaul links for LPGS, as explained in section 5.4.3. However, the coefficients can be set to different values.

The backhaul capacity required at each point is aggregated using the same methodology as for costing up the wireless backhaul in section 5.4.3, with different backhaul capacities aggregated depending on whether the points in the cluster have one unit of demand or more.

This use of a spanning tree means that an LPGS could be linked back to the RAU via one or more BTS and/or other LPGS. Each type of link (LPGS-RAU, LPGS-LPGS, LPGS-BTS, BTS-BTS, BTS-RAU) is identified separately in the output. The number of additional relay stations required between the links are also aggregated.

6 Customer Access Network Design

This section outlines the methodology and calculations used within the CAN module, which is one of the active components of the LRIC model. The inputs for this module have been derived individually for each geotype within the geoanalysis and access network module.

The CAN module contains the calculations that dimension the network assets required from the customer location back to the LE for all customer locations in of Australia.

As illustrated in Figure 6.1 below, the module is comprised of three worksheets containing the inputs and definitions and a fourth worksheet called ‘Access’ that contains the calculations. The ‘Access’ worksheet also provides a summary of assets that are then fed into the Cost module.




Figure 6.1:
Structure of the CAN
module [Source:
Analysys]

The remainder of this section discusses:

- how the CAN is dimensioned
- the separate access technologies and how they are treated in the CAN module
- how the CAN asset volumes for the whole of Australia are calculated within the module
- how a next-generation access network is accommodated within the model.

6.1 Services

As explained in section 3.3, the default input for the number of SIOs which scale the network remains constant at 11.3 million which reflects the potential demand calculated in the location and demand database. This default input can be changed by the user, if a different sized access network is deemed appropriate.

The sample of ESAs used by the geoanalysis and access network module use locations extracted from the Location and Demand Database. These locations also have an associated demand. As a result, assumptions on the demand distribution are embedded within the geoanalysis from this database.

The demand per location parameters derived with the sample of ESAs by the geoanalysis may differ from those calculated using a census of the Location and Demand Database. To be consistent with the other parameters derived from the geoanalysis, the corresponding demand per location parameters are used.

Distribution of total access demand by geotype

The Location and Demand database provides the necessary data to distribute the total access SIOs across the geotypes. This database represents the potential access demand across all of Australia, and therefore the network is dimensioned to reach all locations.

The distribution of total access demand by specific service is further explained in the user guide.

6.2 Treatment of access technologies

The copper CAN is defined as the network that extends from the customer side of the MDF to the first socket in the customer's premises. The equipment side of the MDF is included in the access part of the network.

This definition can be applied for customers served by copper from an LPGA, where the backhaul from an LPGA to the LE would be considered as part of the core network. The allocation of associated costs of the LPGA backhaul is informed by the geoanalysis, where the distance of duct designated as being used just for LPGA backhaul is calculated as a percentage of total duct distance. This distinction of duct is described in more detail in section 5.3.3. The actual cost allocation is performed in the Cost module.

Access services, excluding ULLS, can be delivered over fibre and are recognised in the LRIC model. For fibre, the access network is considered to run from a fibre termination point (FTP) within the customer premises to the LE. At the LE, a fibre network unit (NU) presents the service to the core network in a similar manner to services over the copper network. Therefore, the LRIC model currently considers the FTP, the fibre and underlying duct network and the NU as access elements.

The wireless access technology has been implemented to connect the BTS to the LE through microwave backhaul links, rather than directly into a national mobile network operator. At the LE, it is assumed the service is presented to the core in a similar manner as the copper network. Therefore the LRIC model currently considers the wireless BTS and backhaul as access elements.

Satellite services and the associated annual transponder cost to backhaul to a ground-station in the core network are considered as access elements.

6.3 Calculation of asset volumes in the CAN

The network design in the CAN module is consistent with the architecture described in section 5. The parameters from the geoanalysis and access network algorithms, which inform the CAN module, are average values calculated from the sampled ESAs. These are used to extrapolate the total volumes of assets required for each geotype as a whole. The following section describes the assets modelled.

Exchanges

As discussed in section 4.2.6, a subset of ESAs have been identified to have multiple copper centres. In these cases, the ESA has been subdivided so that each sub-area contains one copper centre, which is recognised as a LE. ESAs identified as only having demand served by satellite are deemed not to have a LE.

The cost of the LE should include a cable entrance facility. The cost of the MDF is calculated separately, based on an average cost of connected SIO. The cost of the MDF should include necessary protection for lightening. Other building equipment, such as air-conditioning and power, to support the platforms are considered through the core module.

Pillars and LPGS

Having calculated the access SIOs within each geotype, the proportion of SIOs served by each of the four access technologies is calculated using proportions calculated from the geoanalysis.

Using further proportions from the geoanalysis, the total copper SIOs in a geotype are further broken down into those served by

- the LE, specifically by a pillar at the LE
- pillars not at the LE
- LPGS.

The number of pillars (excluding those at the LE) required in the geotype is then calculated as the ratio of the copper SIOs served by pillars not at the LE and the average SIOs served by a pillar.

The number of LPGS is derived using an analogous ratio.

Distribution points

The number of distribution points is calculated by the ratio of the total copper SIOs in a geotype to the average SIOs per distribution point.

In urban geotypes, a distribution point can contain up to four locations (one DP and three FDPs), which is an upper bound default assumption used by the DP clustering algorithm. It is important to know the distribution of DPs based on the number of FDPs within the ‘main DP cluster’, as it is used in understanding the network from the DP to the customer location. The ‘main DP cluster’ is the count of locations in a DP cluster which all sit on the main DP-pillar route. A location which is not on this route is known as an ‘isolated FDP’, as shown in

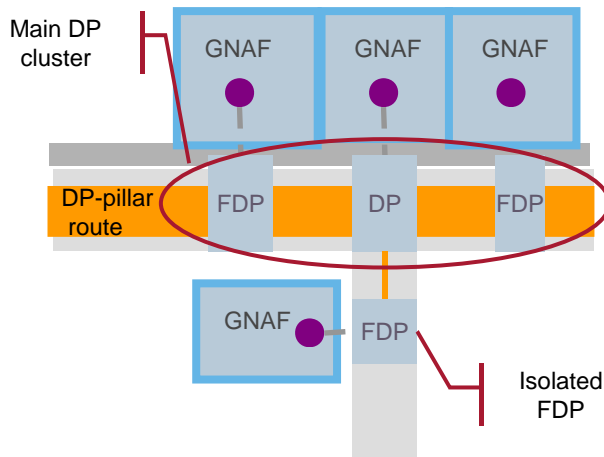


Figure 6.2: Illustration of main cluster and isolated FDP [Source: Analysys]

Fibre rings

The number of fibre rings is derived by calculating the ratio of the total number of nodes that will lie on the rings and the average number of nodes on a ring, which is derived from the geoanalysis. This calculation can also assume that only a proportion of nodes are actually on the fibre ring. This is because the geoanalysis and access network module can also assume that only pillars with existing fibre demand lay on the fibre rings.

Trench network

The trench network from exchange to DP is deployed on one side of the road – the geoanalysis does effectively model the network in the centre of the road, but we believe that this does not cause a material error. We assume an equal probability that locations served will be on the side of the road where the trench is routed or on the opposite side of the road. In reality, there may be occasions where there is some bias of locations to one side of the road the trench would likely be deployed on that side of the road. We have not identified any evidence to suggest this would have a material impact over the whole of the access network.

► Trench network ($PB \gg DP$)

For the architecture for the trench network from a customer location to a DP, we assume that a serving pit (SP1) is deployed at the DP location. In those cases where the DP serves more than 1 location, a road crossing and a second serving pit (SP2) is deployed in certain situations, as shown below.




Figure 6.3:
Structure of a DP cluster
[Source: Analysys]

We therefore need to calculate the trench requirements for the $PB \gg SP$ links and the road crossing ($SP1 \gg SP2$).

We assume that the serving pits are 1m from the PB in all geotypes. The length of the road crossing is then equal to $2 \times (PB \gg FDP) - 2$, which we have calculated using geoanalysis. A road crossing is not deployed for all clusters. Assuming that in a main DP cluster locations are equally likely to be on the side of the road of the DP-Pillar trench or on the other side, a road crossing is not required $1/(2^n)$ times, where n is the number of locations in the main cluster. As an example, where there are two location in a cluster 1 in 4 occasions both of those locations are on the side of the road with the trench.

We assume that the average length of the trench required for the $PB \gg S.P$ links is given by $\sqrt{(x^2 + y^2)}$, where x and y are as shown above in Figure 6.3. We derive x differently depending on four DP cluster types, given by the number of non-isolated FDPs in the main DP cluster. In summary, for DP clusters with:

- 0 non-isolated FDPs, $x=0$
- 1 non-isolated FDP, $x = \text{average distance between locations in DP clusters}^{15}$ for this FDP.

¹⁵ This is calculated in the geoanalysis and access network module, as the average length of FDP-DP trench per FDP

- 2 non-isolated FDPs, x = average distance between locations in DP cluster for both FDPs.
- 3 non-isolated FDPs, x = average distance between locations in DP cluster for two FDPs and $x+x$ for the third FDP.

This is illustrated below in Figure 6.4.

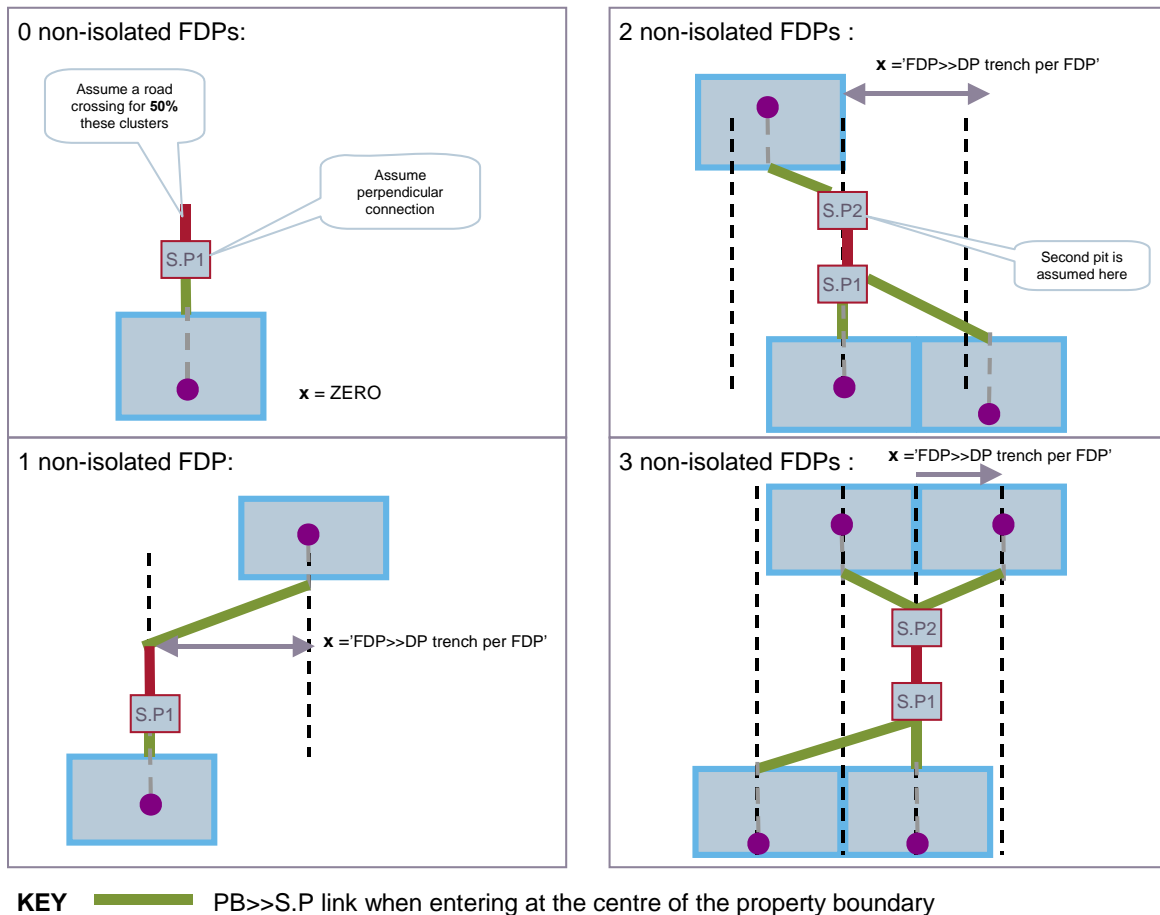


Figure 6.4: Derivation of x by type of DP cluster [Source: Analysys]

We then calculate the average length of the PB>>SP links for each of these four DP cluster types and apply this average length per link to derive the total link lengths.

This assumes that the PB>>SP links connect to the centre of the property boundary. It may be more reasonable to assume that the links enter at the corner of the land parcel boundary. The CAN module allows x to be scaled in order to test this assumption. If we assume that this value is 0.5 for instance, then we find that for DP clusters with:

- 0 non-isolated FDPs, $x=0$
- 1 non-isolated FDP, $x=0.5 \times \text{average distance between locations in DP clusters for this FDP}$.
- 2 non-isolated FDPs, $x=0.5 \times \text{average distance between locations in DP cluster for both FDPs}$.
- 3 non-isolated FDPs, $x=0.5 \times \text{average distance between locations in DP cluster for two FDPs and } (0.5 \times x)+x \text{ for the third FDP}$.

This is illustrated below in Figure 6.5

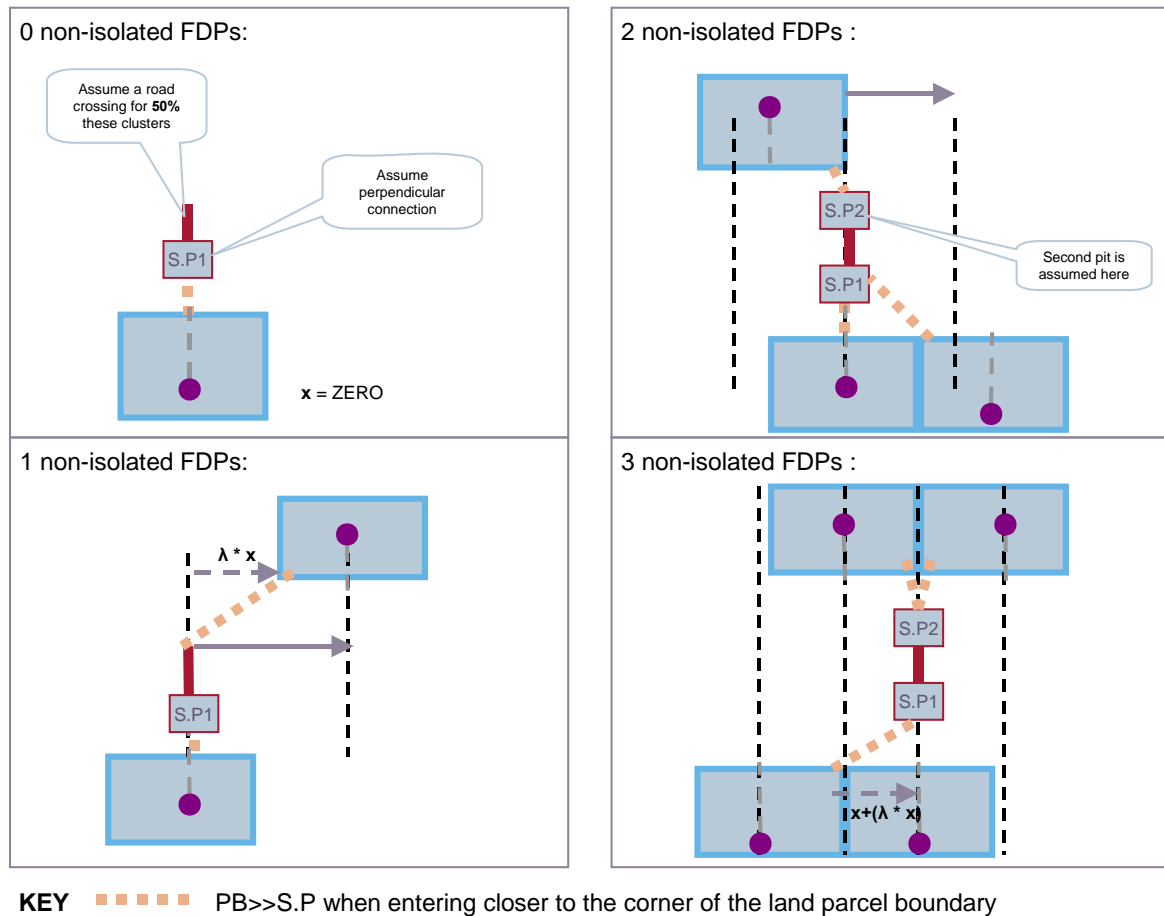


Figure 6.5: Derivation of x by type of DP cluster assuming that the links meet the property boundary closer to the corner of the land parcel [Source: Analysys]

It is a design decision that access to properties is delivered underground, providing a highly robust access network. In reality some access to properties might be delivered through aerial routes from poles or other means. The LRIC model does not currently accommodate such an approach, though adjustments to user inputs may provide an approximate solution.

Trench network ($DP \gg LE$)

The length of trench is calculated for each of the following elements in the path from the DP back to the LE:

- DP >> next node (pillar / LPGS / LE)
- Pillar >> LE
- Links on fibre rings.

These lengths are driven by average trench lengths derived by the geoanalysis.

Having calculated the total trench for each geotype, the duct distributions calculated by the geoanalysis are used to derive the trench length by the number of ducts passing through for all DP-LE links.

Property boundary to serving pit links are assumed to be one duct and are quantified separately so that they are excluded from calculations relating to the transfer of cost between the core and access networks.

Within the cost module, the cost per meter of trench is defined by duct size and includes inputs for the cost of lightning guard wire. Inputs are also provided for duct laid in open trench or for cable ploughed in more rural areas. These inputs are further explained in the user guide.

Pits and manholes

There are five uses of pits that contribute to the total deployment:

- A pit is assumed to be deployed at each distribution point
- Additional P5 pits are also deployed:
 - at any FDP not passed through by the DP-pillar cable deployment (an ‘isolated FDP’)
 - on the remote end of a road-crossing if more than one location is served
- In order to satisfy engineering rules for the maximum permitted haul distance of cable
 - additional pits are included within the trench in the distribution network
 - additional manholes are included within the trench in the main network.

Copper network

► *Copper NTPs*

Network termination points (NTP) at the customer premise are modelled through deploying a 2-pair wall socket or a building termination frame, dependent on the size of the cable used for the lead-in to the location. Building termination frames are modelled in size of 10-pair, 30-pair, 50-pair or 100-pair.

The cost of a 2-pair wall socket is understood to be recovered through the connection charge, therefore the routing factors in the cost module (*RF.Access* sheet) have been set to zero for this asset.

The cost of a termination frame is understood to be associated with the jointing of the lead-in cable to the frame. The frame is understood to be provided by the building owner.

► *Copper sheath*

The length of copper sheath is calculated by cable size, for each of the following elements in the path from a customer location back to the LE:

- NTP >> PB (property boundary)
- PB >> DP (derived using the trench network dimensioned for this part of the path)
- DP >> next node (pillar / LPGS / LE)
- Pillar >> LE.

The links from the pillar / LPGS back to the LE form the *main* copper network, whilst the other elements form the *distribution* network. This distinction is necessary as certain parameters are defined differently between the distribution and main network. Furthermore, this separation supports the deployment of a next-generation access network, described below in section 6.4.

► *Joints and branching kits*

The joints for copper sheath are calculated for the following cases:

- DP >> Pillar (demand) – where a sheath spurs off of the principle distribution sheath. The size of the joint is dependent on the size of the downstream spur cable
- DP >> Pillar (distance) – where the maximum haul distance for the distribution network is exceeded, the whole sheath is jointed
- Pillar >> RAU (distance) – where the maximum haul distance for the main network is exceeded, the whole sheath is jointed.

The number of branch kits is also calculated.

Fibre network

The fibre sheath cable calculations use the same methodology as for the copper sheath cable calculations, with the exception that fibre links cannot be connected from an LPGS directly. Instead, a fibre link within a LPGS cluster is linked all the way back to the exchange. There is also fibre backhaul from the LPGS to the LE.

There is a different calculation for fibre ring distances, where the total length of the ring is dimensioned using the average distance between two fibre pillars and the number of pillar links required.

Wireless network

The number of wireless BTS in a geotype is driven by the proportion of access SIOs designated as wireless and the average size of a wireless cluster. One wireless CPE is deployed per location fed by wireless.

Satellite network

The number of locations served by satellite in a geotype is driven by the proportion of access SIOs designated as satellite and the number of satellite SIOs per satellite location. One satellite CPE is deployed per such location.

6.4 Next-generation access network

A next-generation access network is modelled by the replacement of pillars with MSANs connected to the copper distribution network. This is defined by geotype in the 'Scenario' worksheet of the Cost module.

The MSANs are backhauled on either a fibre ring or a direct fibre link, dependent on whether or not the geotype assumes the deployment of a fibre ring through the pillars. This assumption is defined within the geoanalysis and access network module. The fibre backhaul means that the copper connections from the MSAN back to the LE (i.e. the *main* network) is disregarded. The duct cost supporting the fibre backhaul is informed from the geoanalysis where the duct which facilitated the main copper network is known.

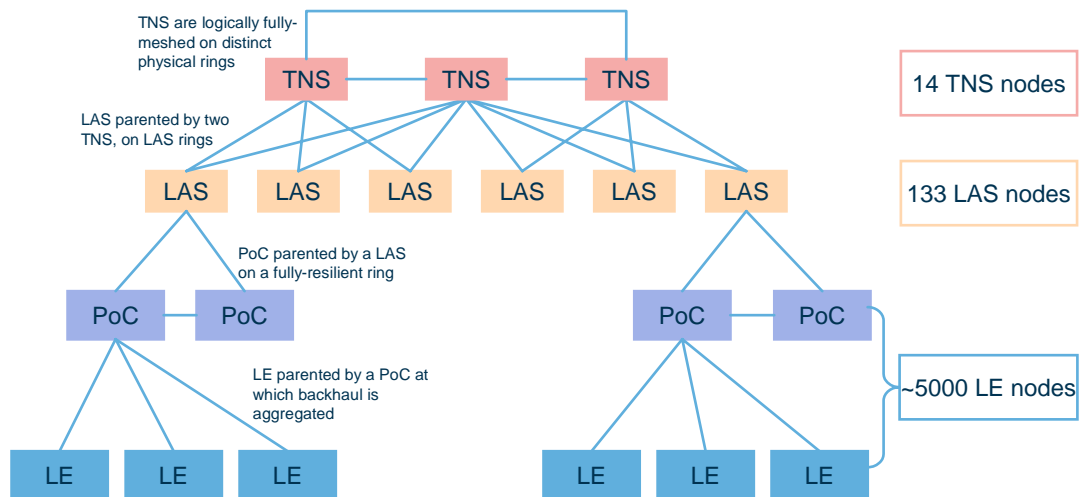
The CAN module has been designed so that the above conversion of assets is possible, however the actual calculations of assets and costs are conducted in the Cost module.

Moving the line card, or NGN equivalent, to positions equivalent to the pillar effectively extends the core network reach. Further discussion on modelling of the NGN is provided below in section 8.

7 Core network design – modern network

This section describes the modern network design used within the Core module and the algorithms implemented to model this design. In implementing both the modern and the next-generation network design, the algorithms developed can accommodate both approaches.¹⁶ Consequently, the worksheet nomenclature takes a generic approach for the two approaches.

The design of the modern core network (see Figure 7.1 below) is based on a TDM architecture in which the voice and data platforms are carried and switched on separate systems, but both carried on the same transmission system (i.e. SDH, plesiochronous digital hierarchy (PDH) or dense wave division multiplexing (DWDM)).



Note: LE: Local exchange; PoC: Point of confluence; LAS: Local access switch; TNS: Transit network switch

Note: A PoC is a local exchange on an SDH ring

Note: Although the Large Pair Gain Systems (LPGS), such as CMUX equipment, are costed as part of the core network, the deployment of these assets is actually calculated in the geoanalysis and access network module as it is modelled as an access decision.

Figure 7.1: Design of the modern core network [Source: Analysys]

The LAS and TNS are located, respectively, at LTH and MTH. The data platforms – which are responsible for switching and routing the data traffic – are also co-located at the LTH and MTH sites.

Pair gain system requirements are calculated in the CAN module, however, they are identified as a core network asset (together with the fibre links from the pair gain system back to the LE). These assets are consequently allocated to the core service costing in the Cost module.

¹⁶ A modern network design is implemented by setting all of the 'MSANs deployed in geotype' parameters to 0. Defined in 'Geotype.served.by.MSAN.equipment' named range on the 'Scenario' sheet of the Cost module, passed to 'In.Control' sheet of the Core module.

The core network equipment for the NGN and/or data networks (broadband remote access servers (BRAS), user authentication servers (RADIUS), domain name servers (DNS), core network management systems, intelligent network (IN) platforms, billing equipment) is located at the MTH location.

The modelling of the core network follows a scorched-node principle, maintaining the location of network nodes as they exist in the current network, i.e. 14 TNS node locations, 133 LAS node locations and up to 5254 LE node locations.¹⁷

The required output from the Core module is the total number of assets deployed. The total number of assets is dimensioned by the traffic carried over them. The usage of certain assets, such as trenches and ducts, is explicitly calculated as part of the network deployment algorithms – the resulting usage intensities (the routing factors) are employed to allocate the cost of these network assets (in terms of minutes or kbit/s conveyed on the network) between the various platforms that use them. The total output (number of assets) from the Core module is on the ‘Out.Assets’ worksheet.

The model uses cost-minimising backhaul algorithms in order to determine a cost-efficient backhaul architecture at the LE and PoC level, whilst maintaining realistic network planning principles such as the maintenance of a resilient¹⁸ network structure.

Deployment of core duct and fibre in areas where an access network exists utilises the existing trench of the access network. This results in a small cost saving which is shared between the CAN and the core network.

Figure 7.2 below illustrates the calculation flow of the Core model.

¹⁷ Analysys has identified 5070 ESAs and 5117 copper centres; individual ESA-copper centres combinations total 5254.

¹⁸ A network is considered resilient if the ring is cut at any one point, but traffic will still be able to be automatically routed.

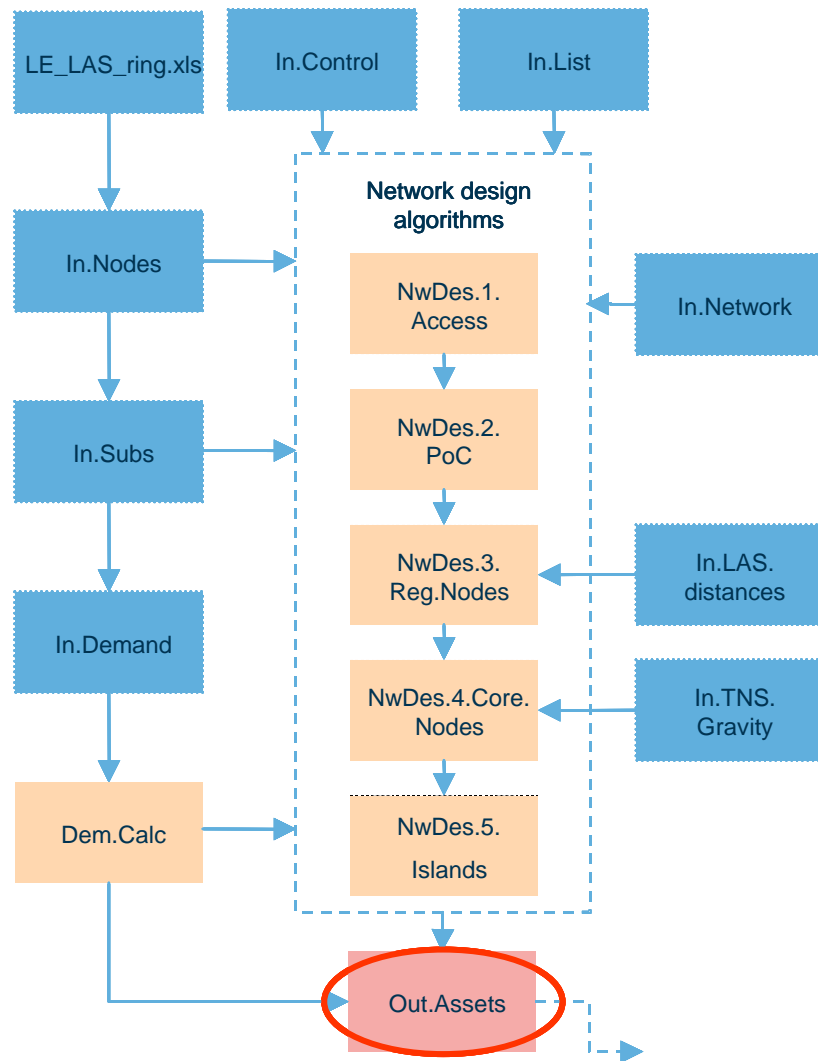


Figure 7.2: Calculation flow for the Core module [Source: Analysys]

Input data is entered into the Core module on the worksheets prefixed ‘In.’ Scenario control data is linked into the Core module from the Cost module on the ‘In.Control’ worksheet.

Commonly used named ranges are listed on the ‘In.List’ worksheet.

Node data for each level in the core network (LE, PoC, LAS and TNS) are listed on the ‘In.Nodes’ worksheet. Certain data (LE-PoC and PoC-LAS data) is linked in from the *Le_LAS_ring.xls* external workbook.

Subscriber numbers for each level in the core network (LE, LAS and TNS) are calculated on the ‘In.Subs’ worksheet. This data is used in the network design algorithm calculations.

Demand data is linked into the Core module from the Cost module on the ‘In.Demand’ worksheet. This data is used in the network loading calculations on the ‘Dem.Calc’ worksheet.

Network parameters which are used in the calculations in the network design algorithms are input on the 'In.Network' worksheet.

Information on the distances between LAS nodes is contained in the 'In.LAS.distances'. This information is specifically used on the 'NwDes.3.Reg.Nodes' worksheet.

Information on the traffic levels on the transit network is calculated on the 'In.TNS.gravity'. This information is specifically used on the 'NwDes.4.Core.Nodes' worksheet.

The output data (number of assets required, network link usage between platforms) from each of the network design algorithm worksheets is collated on the 'Out.Assets' worksheet. This data is subsequently linked into the Cost module.

The remainder of this section discusses:

- demand drivers used to dimension the core network
- the principle of using routing factors to mapping service demand onto the loading of the network elements
- the principle of quality of service in the core network calculations
- modelling of the transmission links in the modern network
- dimensioning of equipment and transmission at the LE level
- dimensioning of the equipment and transmission at the PoC level
- dimensioning of the equipment and transmission at the LAS level
- dimensioning of the equipment and transmission at the TNS level
- dimensioning of special backhaul solutions required for customers located on islands
- how the overlap between access and core networks is evaluated.

7.1 Demand drivers

The main demand drivers used to dimension the various network elements in the core network are discussed in this section. These demand drivers are derived from the total service volumes that are contained within the Cost module (refer to section 3).

The demand loading on the modern core network is calculated on the 'Dem.Calc' worksheet.

The following subsections discuss:

- the calculation of the busy-hour voice traffic
 - including the additional calculation for the inclusion of ringing time
- the calculation of the busy-hour ISDN traffic
- the calculation of the average bandwidth provisioning for broadband services
- the inclusion of other transmission services in the core network.

Voice services (Sheet Dem.Calc, Cells E59:O88)

In the modern network, the number of E1s required to carry the network traffic needs to be dimensioned. In order to do so, the number of voice minutes is converted into a year-average busy-hour Erlang (BHE) load using the following inputs:

- proportion of annual traffic during 250 normal¹⁹ weekdays
- proportion of weekday traffic occurring in the normal busy hour
- the average proportion of daily calls that occur in the busy hour.

$$BHE = \text{annualtraffic} \times \frac{P_d \times P_w}{B_d \times 60}$$

Where:

P_d = Proportion of daily traffic in the busy hour

P_w = Proportion of annual traffic in the busy week days

B_d = Number of busy (week) days

Figure 7.3:
Calculation of the
busy-hour Erlang
voice demand
[Source: Analysys]

The number of voice BHE is converted into a further measure of demand, namely the number of busy-hour call attempts (BHCA) using the following inputs:

- average call duration
- number of call attempts per successful call (e.g. due to unanswered calls).

$$BHCA = \frac{BHE \times C}{D_{ave}}$$

Where:

C = Call attempts per successful call

D_{ave} = Average duration of a successful call

Figure 7.4:
Calculation of the
number of busy-hour
call attempts
[Source: Analysys]

Ringling time

Voice services explicitly include the additional Erlang load presented by the ringing time associated with calling. Ringling time occurs for calls to a land line where there is network occupancy until the call is answered, diverted or not answered. A ringing time of 10 seconds for answered calls and 28 seconds for unanswered calls to an end user is applied to the various call types.

For each service, the Core module calculates the occupancy minutes in the network:

¹⁹ Normal being defined as a day which is not a bank holiday.

$$\text{Occupancy minutes} = C \times (D_{ave} + R_{suc}) + (CA - C) \times R_{unsuc}$$

Where:

C = Successful calls

D_{ave} = Average duration of a successful call

R_{suc} = Average ringing and call set-up time for successful calls

CA = Total call attempts (successful calls + unsuccessful calls)

R_{unsuc} = Average ringing and call set-up time for unsuccessful calls

Figure 7.5:
Calculation of the
total occupancy
minutes [Source:
Analysys]

ISDN services (Sheet Dem.Calc, Row 69)

ISDN services have been principally included in the Core module as the line cards can share the same racks as PSTN line cards, and consequently local exchange accommodation costs. In addition, ISDN has some impact on backhaul requirements, though of a declining significance relative to voice and more recently xDSL services.

ISDN impact on the core network is calculated in BHE and sequentially E1s in the same manner as voice, described above.

Demand for broadband services (Sheet In.Network, Rows 44:48)

ADSL and SDSL services are modelled as being delivered through an ADSL2+ capable DSLAM, which accommodates line cards capable of ADSL2+ or SDSL business services. The backhaul from an exchange is provisioned on the existing SDH network thus creating economies of scope. The capacity is provisioned by defining a number of E1 equivalents for each xDSL rack installed. As shown in Figure 7.6 below, a DSLAM of 48 ports per card, 18 cards per shelf and 4 shelves per rack gives a maximum capacity of 3456 ports. Provisioning an E3 (14 E1s) gives an effective payload capacity of 27Mbit/s (14 x 1920kbit/s²⁰) or minimum of 8kbit/s per subscriber assuming full port utilisation.

ADSL2+ specification	Maximum unit capacity	Backhaul provisioned	Effective backhaul provisioned
Ports per line card	48	-	-
Line cards per shelf	18	-	-
Ports per shelf	864	E3 = 16 x E1 (31Mbit/s payload)	36kbit/s
Shelves per rack	4	-	-

Figure 7.6:
Parameters used in
the calculation of the
backhaul provisioning
of xDSL SIOs
[Source: Analysys]

²⁰ Based on 30x64kbit channels, with two channels reserved for signalling.

Over the whole DSL network, the average backhaul per subscriber presented at the LTH can be determined and is used to dimension core network links.

Other transmission services (Sheet Dem.Calc, E82:E84)

Transmission volumes associated with other uses of transmission capacity other than PSTN and xDSL traffic are explicitly modelled in the core network. 'Other transmission' represents traffic from sources such as the declared transmission and digital data access service (DDAS), leased lines, traffic from mobile networks, cable-TV content distribution, etc. It is modelled as a Mbit/s load which dimensions the links in the core network.

7.2 Routeing factors (*Sheet Dem.Calc, Rows 91-270*)

An input table of routeing factors determines the factor applied to each service volume when calculating the load on the various parts of the network.

The routeing factors for a particular traffic service are calculated on the basis of the number of times/loading that a particular network element is used to deliver the service being modelled. Different combinations of network elements may be used depending on the path taken in the network. For example, PSTN local traffic may be switched by only one LAS, or may be switched by two LAS, or may indeed involve switching at the transit layer.

The proportion of calls that utilise a particular combination of network elements is used to ascertain the average routeing factors for that particular type of traffic.

7.3 Quality of service (*Sheet NwDes.1.Access, H10533:H15792*)

The Core module explicitly takes into account a quality-of-service (QoS) parameter in the network by means of a network-blocking probability. This parameter represents the end-to-end probability that a call is blocked due to network congestion on a link(s) required to complete the call. The module assumes an end-to-end network blocking probability of 0.5% (grade of service (GoS)).

This factor is taken into account in terms of the Erlangs-to-channel conversion (i.e. the number of channels required to provide capacity for a defined number of conveyed Erlangs).

The Core module converts the BHE load into an Erlang channel requirement using the Excel *NORMINV* function, which approximates the Erlang-B formula:

- *NORMINV* ($p, mu, sigma$) returns the value x such that, with probability p , a normal random variable with mean mu and standard deviation $sigma$ takes on a value less than, or equal to, x .

$$p = (1-GoS)$$

$$\mu = \text{Traffic in BHE}$$

$$\sigma = \text{ErlangConversionFactor} \times \sqrt{BHE}$$

Figure 7.7:
Parameters used in
the calculation of the
Erlang-B formula
[Source: Analysys]

- The result of the NORMINV formula is divided by the number of circuits in an E1 to calculate the number of E1s required. A rounded-up number is calculated for each traffic type (PSTN/ISDN/xDSL).

The Erlang formula is non-linear at low numbers of channels, however, it becomes broadly linear in nature at higher channel usage, as seen at local exchanges and higher-level network elements. Consequently, the Core module employs an Erlang conversion factor which maps the channel-circuit relationship at a high channel number.²¹

7.4 Transmission modelling

In the TDM modern network design, the transmission network is dimensioned in terms of E1 virtual containers (VC).

Unlike in the NGN IP network, only transmission of the same type may be aggregated, i.e., PSTN, ISDN and xDSL traffic is maintained separately. Thus, the Core module calculates the number of E1 VCs required to handle the PSTN, ISDN and xDSL traffic separately.

The capacity of synchronous digital hierarchy (SDH) equipment (STM-x²²) in terms of E1 VCs is known from industry standards:

PDH/SDH transmission level	Number of E1 VCs
E1	1
E3	14
STM-0	21
STM-1	63
STM-4	252
STM-16	1008
STM-64	4032

Figure 7.8: PDH/SDH
transmission –
capacity of E1 VCs
[Source: Analysys]

The Core module considers the number of units of a single transmission speed needed to be deployed to meet the number of E1 VCs required. This calculation also takes into account the cost threshold associated with deploying these transmission assets. It is assumed, based on benchmark

²¹ Use of the proper Erlang B formula with a look-up function results in a significant slowing of the model recalculation time.

²² Synchronous Transport Module (STM) xth level frame structure

data, that such cost increases approximately 2.5 times with respect to a quadrupling of speed, i.e. an STM-4 is approximately 2.5 times more expensive than an STM-1. Consequently, instead of deploying three STM-1 links, the Core module will deploy a cheaper solution of one STM-4 link.

7.5 Dimensioning at the LE level

The 'NwDes.1.Access' worksheet contains the calculations for the dimensioning of equipment and transmission at the LE level. Calculations are performed individually for each LE. This approach enables the accurate dimensioning of traffic flows at the LE level, and at aggregated levels higher in the network. A summary of the modern core network units deployed in a year at the LE level are provided at the bottom of the worksheet (*Sheet NwDes.1.Access, Rows 21073-21096*).

7.5.1 Dimensioning of LE equipment (*Sheet NwDes.1.Access, Rows 7-5268*)

The MDF at the LE is defined as the physical boundary between access and core networks. In the modern network, the LE contains a digital remote concentrator²³, which multiplexes voice circuits back to the LAS. Remote concentrator units are scalable (in that they are composed of shelves and racks) and are dimensioned according to the number of required line cards.

In the modern network, separate equipment platforms are required for voice (PSTN and ISDN) and data (ADSL and SDSL) traffic. For the former, Ericsson AXE equipment is modelled, whilst the latter is modelled using a DSLAM which can utilise the SDH network in place, such as Alcatel 1540 equipment:

- **Ericsson AXE** – the configuration of this equipment is assumed to be a maximum of 48 ports per PSTN line card; 24 ISDN-BRI ports and 16 ISDN-PRI ports per ISDN line card. PSTN and ISDN (BRI and PRI) ports are modelled to be deployed on separate line cards. Each shelf can hold a maximum of 16 line cards, with a maximum of 3 shelves per chassis.
- **Alcatel 1540** – the configuration of this equipment is assumed to be a maximum of 48 ports per xDSL line card, with a maximum of 18 line cards per shelf. There is a maximum of 4 shelves per chassis.

²³ The remote concentrator is distinct from a C-MUX (customer multiplexer) which may be deployed outside of an exchange for in-fill or long-reach needs.

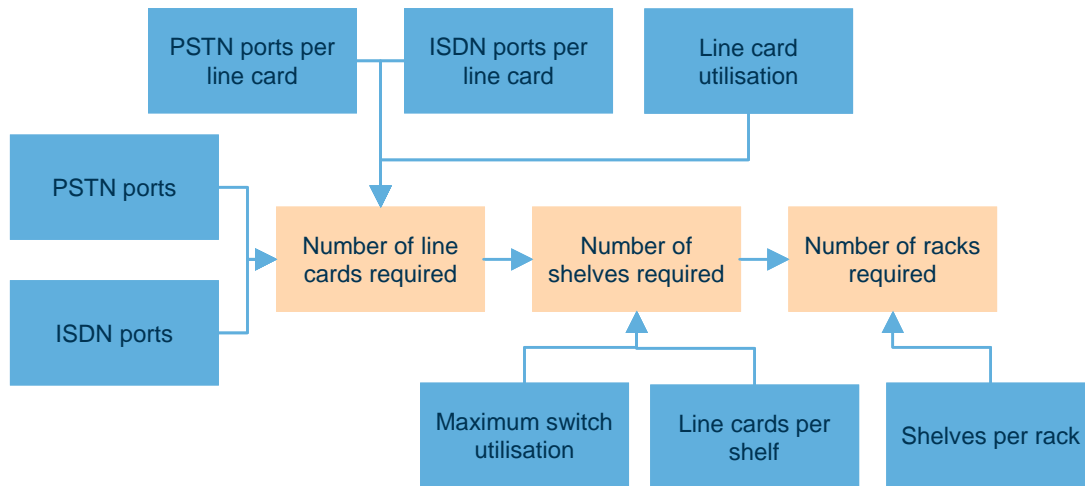


Figure 7.9: Calculation of equipment platform units and line cards required [Source: Analysys]

Utilisation parameters are defined in terms of line cards, shelves and racks, which allow for some realistic over-provisioning of assets at each local exchange. Current assumptions are shown in Figure 7.10.

Network element (capacity measure of element)	Voice	xDSL
Line card (per port)	100%	100%
Shelf (per line card)	90%	90%
Rack (per shelf)	100%	100%

Figure 7.10:
Maximum utilisation
parameter of local
exchange assets
(Sheet In.Network
Rows 57:60)
[Source: Analysys]

It is appropriate to note that given each exchange is explicitly modelled, it is appropriate to minimise over-provisioning compared to a entire network model which takes a higher level perspective.

7.6 LE transmission architecture

In terms of the LE backhaul transmission architecture, a number of LEs are connected to a point of confluence (PoC) – which itself is an LE location.

The POC is an LE which is large enough to be on a protected SDH Ring and from which a number of smaller exchanges on unprotected transmission capacity connect into the SDH ring.

The efficient transmission route between the LEs and the parent PoC is determined by using a minimum spanning tree algorithm that uses a cost-minimising function. This function reduces the cost of deployment based on trench and fibre costs. The minimum spanning tree code is outlined in detail in the document *Description of the Visual Basic in the Fixed LRIC model*.

The implementation of this network architecture enables demand from each of the LEs to be aggregated at the PoC level in order to backhaul the traffic to the LAS. This represents a more cost-efficient backhaul methodology rather than having individual links back from each LE to the LAS.

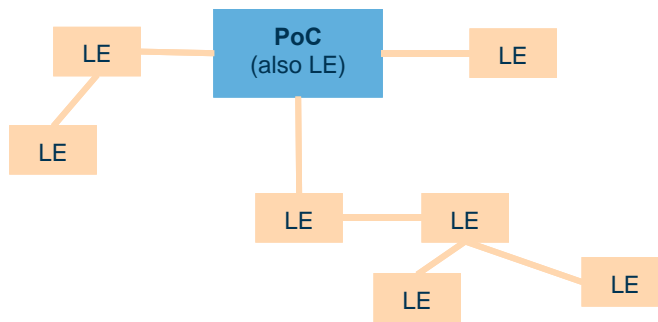


Figure 7.11:
Minimum spanning
tree representation
of LE-PoC
transmission
[Source: Analysys]

The PoC architecture also enables resilience to be implemented in the network (this principle is outlined in section 7.7).

The principles employed in the modelling of the LE-to-PoC links are as follows:

- an LE that serves more than 3000 SIOs is designated as a PoC²⁴
- the remaining LEs are clustered together according to the following dimensioning rules:
 - the number of LEs connected to a PoC is limited to 16. This reflects the objective that no more than 3000 customers should be isolated by a single cable cut
- LEs with the same parent PoC are linked together using the minimum spanning tree algorithm (outlined in Annex C.2)

7.7 Dimensioning at the point-of-confluence (aggregation node) level

The *NwDes.2.PoC* worksheet contains the calculations for the dimensioning of equipment and transmission at the point-of-confluence (PoC) level. A PoC is a LE at which traffic is aggregated in order to backhaul the conveyed traffic to the LAS via a resilient ring architecture. A summary of the total number of units of modern network equipment used in the core network at the PoC level is provided at the bottom of the worksheet (*Sheet NwDes.2.PoC, Rows 2035-2079*).

7.7.1 PoC transmission architecture

A number of PoCs are connected to the parent LAS by means of SDH rings. This ring structure ensures resilience of traffic flows, so that if the ring is cut at any one point, then the traffic will still be able to be automatically routed. This is illustrated in Figure 7.12 below.

²⁴

Note: this is in alignment with the statement in the PIE II documentation.




Figure 7.12:
PoC-LAS network
topology [Source:
Analysys]

Note: In this example, three rings map directly onto the LAS node, with a child ring (coloured in purple) linking directly onto a parent ring (coloured in light blue) via two bridging nodes.

The structure in Figure 7.12 includes the following functionality:

- The algorithm seeks to minimise the distance-based cost associated with the deployment of trench and fibre.
- It further seeks to minimise cost through the implementation of multi-rings, in which a child ring is linked to the parent ring by means of two bridging nodes for resilience purposes.

The dimensioning of these rings is calculated using a multi-ring travelling salesman algorithm.

7.7.2 Multi-ring travelling salesman problem algorithm

The principles employed in the modelling of the PoC-to-LAS links are as follows:

- PoCs may be connected to the parent LAS using optimised multiple SDH rings
- a particular ring may not contain more than 8 PoCs, with a maximum of 8 links between a PoC and its parent LAS
- where multiple rings are modelled, the capacity of a child ring (a ring which does not contain the parent LAS) must also be carried on the parent ring (the ring which contains the parent LAS)
- to ensure resilience, a child ring must be connected to the parent ring by means of two bridging nodes. Where a child ring connects to the parent ring at the LAS node, only a single bridging node is required
- an add-drop multiplexer unit (ADM) is required at each PoC on a ring. The capacity of the ADM is dimensioned according to the size of the total traffic carried on the SDH ring.

The inputs for the LE–PoC transmission are calculated in an external workbook,²⁵ and are copied into the Core module as pasted values into the ‘In.Nodes’ worksheet, specifically *Sheet In.Nodes, Rows 5267-6766*.

The algorithm used to generate the multi-ring structures is described in section 3 of *Description of the Visual Basic used in the fixed LRIC model*.

7.8 Dimensioning at the LAS level

The ‘NwDes.3.Reg.Node’ worksheet contains the calculations for the dimensioning of equipment and transmission at the local access switch (LAS) level. As per the scorched-node principle, the LAS calculations are performed for each of the 133 LAS locations currently in the network.

The LAS itself is responsible for the routing and switching of voice traffic and, therefore, provides the PSTN service features. It also aggregates traffic from the LE levels, returns all local calls and passes on the trunk calls either to other LAS or to the TNS.

A summary of the total number of units of modern network equipment used in the core network at the LAS level is provided at the bottom of the worksheet (*Sheet NwDes.3.Reg, Rows 1969-2036*).

7.8.1 Dimensioning of LAS equipment

The LAS is modelled to be located at a LTH – this reflects the current network design. It is assumed that data routing and switching is performed at the main transit hub site, i.e. it is assumed that the data network has no capability to route at the LAS/LTH level.

Building-related equipment

The LRIC model calculates the cost of building-related equipment, including:

- **generator** – used to supply electricity to all the equipment housed at the LTH
- **uninterruptible power supply (UPS)** – a power supply that includes a battery to maintain power in the event of a power outage
- **air-conditioning unit** – used to maintain the environment required in the LTH.

²⁵ This external workbook contains the graphical results for all of the LAS-area multi-ring networks.

Voice network elements (Sheet NwDes.3.Reg.Nodes, Rows 150-290)

The LAS consists of a switchblock and processor element – these are dimensioned separately in the Core module. The switchblock is dimensioned on the basis of BHE traffic handled by the switch:




Figure 7.13:
Dimensioning of LAS
switchblock units
[Source: Analysys]

The processor is dimensioned according to the processing load, measured in BHCA, handled at each LAS site.

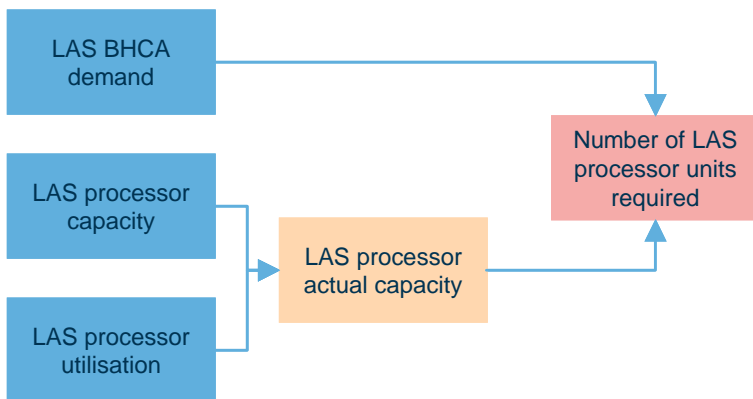


Figure 7.14:
Dimensioning of LAS
processor units
[Source: Analysys]

The number of LAS chassis required is calculated as the maximum number of switchblocks and processors needed at each LAS site.

An ISDN platform is assumed at each LAS, but not explicitly dimensioned as it is not relevant to the declared services. However, the BH traffic demand is of some relevance to dimensioning the transmission links.

Data network elements


It is assumed that an ATM switch is not required at the LTH location as there is no local routing of data network traffic.

7.8.2 LAS transmission architecture

The Core module implements a LAS-ring transmission architecture – a series of rings are deployed which provide a fully-resilient backhaul link between each of the LAS units and a pair of parent TNS units. These rings are individually dimensioned by LAS-LAS traffic and LAS-TNS traffic carried on them.

The principles employed in the modelling of the LAS rings are as follows:

- each LAS ring contains two TNSs
- traffic flowing between LAS sites that are both co-located with TNS sites does not dimension the LAS rings – this traffic is carried on the transit rings. This is conceptually shown in the diagram below:



The Core module contains the flexibility to accommodate different set-ups for the LAS rings.²⁶

The present ring structures are based on Analysys’s interpretation of an efficient manner in which to link together the LAS nodes.


It is recognised that the optical fibres are mostly deployed along roadways (and in some cases along railway lines). Whilst the Core module is currently set up with straight-line distances for LAS-LAS links (identified on the ‘In.LAS.distances’ worksheet), an uplift is applied to correct for road-length distances. The uplift is based upon the overlap analysis, described below in 7.11.6.

It is further recognised that a realistic deployment would be the least-cost pathway linking LAS nodes via a number of exchanges, however, modelling such a deployment is extremely difficult. As a proxy for this, the model assumes sharing of the trench and duct between the different transmission levels in the network, again described below in 7.11.6. This mitigates the costs incurred. Fibre regenerators are deployed for every 80 kilometres of fibre sheath²⁷.

²⁶ The model user can overwrite the present LAS ring nodes with an alternative set. See section 5.12 of the user guide for further details.

²⁷ It is acknowledged that the regenerator would, where possible, be located at a pre-existing node site (such as a local exchange building) – in which case an ADM would be used to drop off capacity at that exchange location. However, the model does not explicitly take this into account. By following roads, it is likely that routes pass local exchange buildings where a regenerator would be deployed

The major LAS rings modelled nationally is shown in Figure 7.16. Rings in the main metropolitan areas in Australia are outlined in Figure 7.17 to Figure 7.22 below.



Note: LAS locations are indicated by the blue stars; red lines indicate the routing of the LAS ring structures. The lines drawn are straight-lines between the appropriate LAS nodes on the LAS rings. In the model, an uplift is applied on these straight-line distances to account for the non-straight-line nature of the road/railway distances associated with the LAS ring links.

Figure 7.16: LAS ring structures – national level [Source: Analysys]




Figure 7.17:
LAS ring structures –
Perth [Source:
Analysys]




Figure 7.18:
LAS ring structures –
Adelaide [Source:
Analysys]

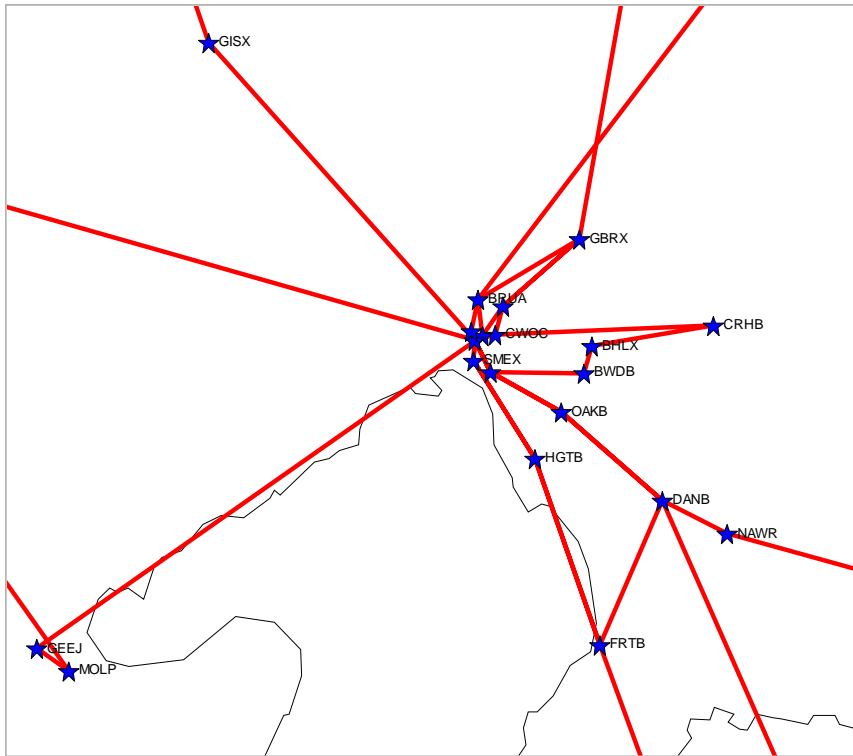


Figure 7.19:
LAS ring structures –
Melbourne [Source:
Analysys]




Figure 7.20:
LAS ring structures –
Canberra [Source:
Analysys]

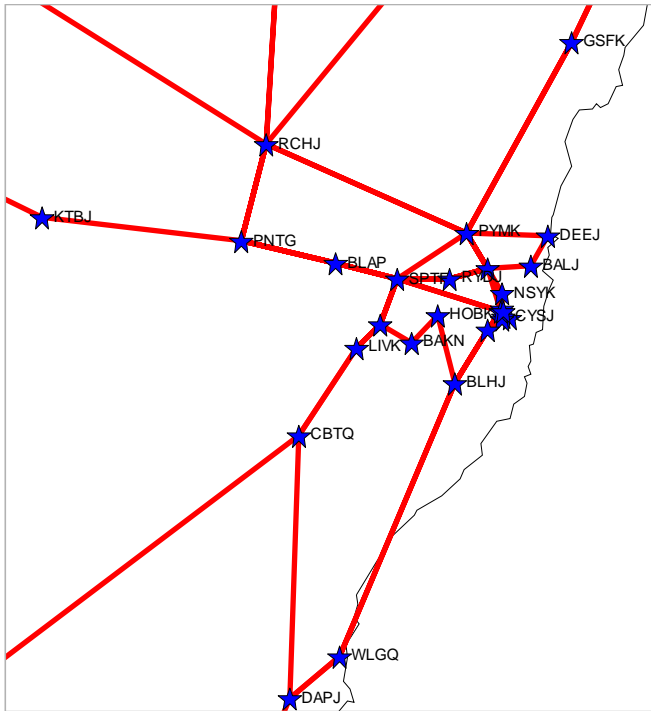


Figure 7.21:
LAS ring structures –
Sydney [Source:
Analysys]

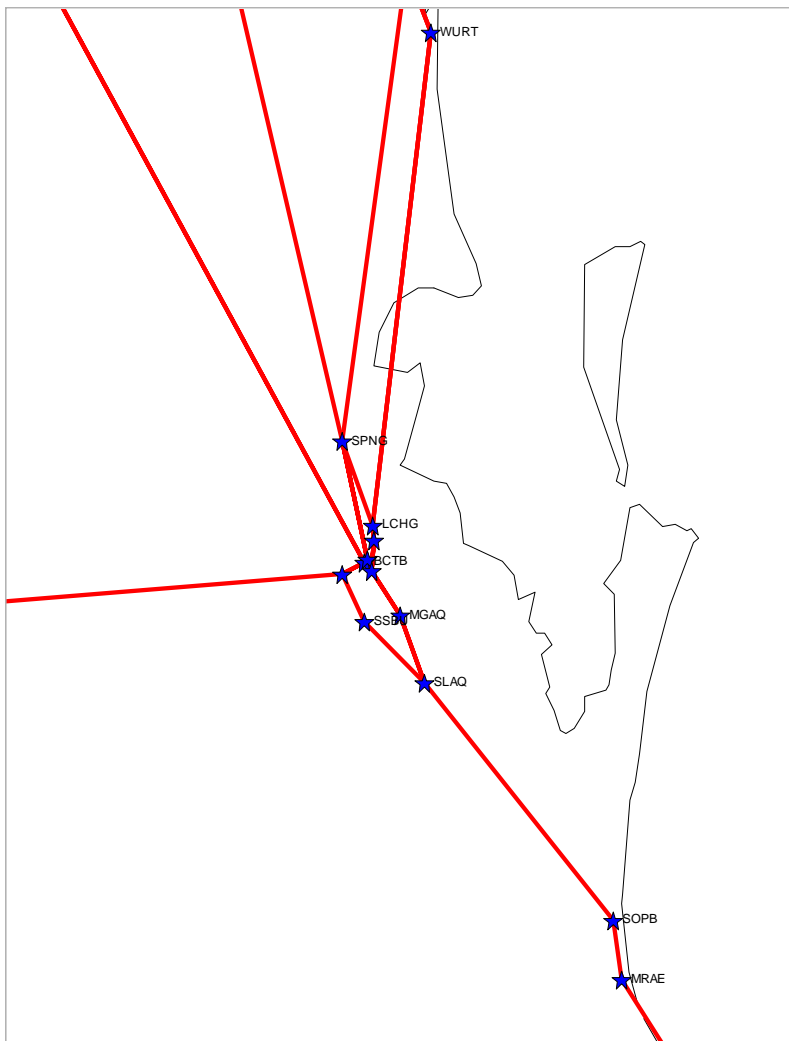


Figure 7.22:
LAS ring structures –
Brisbane [Source:
Analysys]

The Core module can deploy dense wavelength division multiplexing (DWDM) transmission equipment for LAS rings. This parameter is controlled manually within the Core module with a selectable switch on the 'In.Control' worksheet. Deployment of DWDM links is dependent on the ring requiring n or more STM-64 links, where $n=2$ in the base case implementation. It should be noted that it is assumed that the jointing costs for fibre are included in the equipment cost.

The Core module also accounts for traffic interconnected with other networks at the LAS location. The level of interconnection traffic is calculated on the 'Dem.Calc' worksheet. Interconnection is assumed to be delivered through STM-1 ports, calculated based on the number of E1 equivalent ports required.

7.9 Dimensioning at the TNS level

The 'NwDes.4.Core.Nodes' worksheet contains the calculations for the dimensioning of equipment and transmission at the TNS level. As per the scorched-node principle, the TNS calculations are performed for each of the 14 TNS locations currently in the network.

It is modelled that dial-up traffic is carried to the LTH at which point it is switched on to the data network. This dial-up traffic is subsequently carried to the MTH before it is interconnected to the Internet gateway switch. Broadband traffic is carried on the data network and is also interconnected to the Internet gateway switch from the MTH location.

The volume of traffic routed across TNS-TNS links is determined by the routing table parameters. Mobile PoI are modelled to be primarily located at the TNS in each capital city, though a proportion of traffic can be handed over at the LAS level using parameters on the 'Dem.Calc' sheet. This interconnection process occurs through a gateway switch. The Core module further assumes that international gateway switches are available at each capital city – it is acknowledged that this differs to the current network in which international gateways are understood to be only located in Sydney and Perth.

The handover of traffic at the nearest available point to the end-user (termed 'near-end handover') is assumed where possible. However, this can only occur if the alternative carrier has the ability to receive calls at that PoI, otherwise it is handed over at the nearest available PoI.

A summary of the total number of units of modern network equipment used in the core network at the TNS level is provided at the bottom of the worksheet (*Sheet NwDes.4.Core.Nodes Rows 555-633*).

7.9.1 Dimensioning of TNS equipment

The TNS is modelled to be located at an MTH. The TNS itself is responsible for the routing and switching of voice traffic, however, the data routing and switching equipment is also located at the MTH.

Building equipment

The Core module costs the equipment associated with the operation of the building equipment, which includes the following elements:

- **generator** – used to supply electricity to all the equipment housed at the MTH
- **uninterruptible power supply (UPS)** – a power supply that includes a battery to maintain power in the event of a power outage
- **air-conditioning unit** – used to maintain the environment required in the MTH.

Voice network elements (Sheet NwDes.4.Core.Nodes Rows 95-109)

The TNS consists of a switchblock and processor element – these are dimensioned separately in the Core module. The switchblock is dimensioned on the basis of the BHE traffic handed by the switch, whilst the processor is dimensioned according to the processing load, measured in BHCA, handled at each TNS site. The number of TNS chassis required is calculated as the maximum number of switchblocks and processors needed at each TNS site. This calculation follows the same methodology as for the LAS equipment.

An ISDN platform is assumed at each TNS, but not explicitly dimensioned as it is not relevant to the declared services. However, the BH traffic demand is of some relevance to dimensioning the transmission links.

Data network elements

The MTH contains an ATM platform and routers which are used to switch and route the data traffic.

Other network elements

Other core network equipment is located at the MTH location, including:

- **Network synchronisation equipment** – many services running on modern digital telecoms networks require accurate synchronisation for their correct operation. For example, if switches do not operate with the same rate clocks, then slips will occur degrading performance. Telecoms networks rely on the use of highly accurate primary reference clocks which are distributed across the network using synchronisation links and synchronisation supply units (SSU):
 - **primary reference clock** – these highly accurate units must meet the international-standards requirement for long-term frequency accuracy better than 1 part in 10^{11} . Such systems are usually based on atomic clocks or receivers using GPS

- **SSU equipment (plus licence)** – this equipment is used to ensure reliable synchronisation distribution.
- **STP** – this is a router that relays SS7 messages between signalling end-points (SEP). The STP is connected to adjacent SEPs and STPs via signalling links. Based on the address fields of the SS7 messages, the STP routes the messages to the appropriate outgoing signalling link. In order to meet resilience requirements, STPs are modelled to be deployed in mated pairs at each node.
- **Core network management system (NMS)** – this equipment is a combination of hardware and software used to monitor and administer a network.²⁸
- **Intelligent network (IN)** – it allows operators to provide value-added telephony services, such as televoting, call screening, telephone number portability, toll-free calls, prepaid calling, account card calling, virtual private networks, etc. The cost of this equipment should not be apportioned to termination costs, as it is recovered through charges for value-added services (VAS).

Specific equipment is required at the MTH for DSL traffic:

- **Broadband remote access server (BRAS)** – it routes traffic to and from the DSLAMs located at the local exchange. The BRAS is also the interface to authentication, authorisation and accounting systems (such as RADIUS). The number of BRAS required is calculated as the total number of DSL subscribers logged on at any one time divided by the capacity of a BRAS.
- **Domain name server (DNS)** – this equipment translates human-readable computer hostnames, e.g. www.example.com, into the IP addresses, e.g. 208.77.188.166, to which networking equipment needs to deliver information. It also stores other information such as the list of mail exchange servers that accept email for a given domain.
- **Remote authentication dial-in user service (RADIUS) (and accounting system)** – it is an authentication, authorisation and accounting (AAA) protocol for controlling access to network resources. It is used to manage access to the Internet or internal networks across an array of access technologies including modem, DSL, wireless and VPNs.

Specific to dial-up Internet services, a dial-up remote access server (RAS) is included at each MTH node to route dial-up traffic to the Internet. The RAS is also the interface to authentication, authorisation and accounting systems (such as RADIUS).

7.9.2 TNS transmission architecture

Transit network switches are logically fully meshed in a series of physical ring structures.

²⁸

Note: the NMS that is modelled at each TNS location, is assumed to include the cost of the centralised network management centre.

7.9.3 TNS logical ring dimensioning

The traffic flowing across the transit network is dimensioned according to the outputs of the routed demand on a per-subscriber basis. The number of subscribers (PSTN, ISDN, xDSL) at each TNS is calculated according to the logical dimensioning of the LAS–TNS links. The destination of this traffic from each TNS is modelled using a gravity model – the calculations for which are located on the ‘Input.TNS.gravity’ worksheet. This based on the following formula:

$$\frac{P_1 * P_2}{d^k}$$

Where:

P_1 = *Subscribers at TNS1*

P_2 = *Subscribers at TNS2*

D = *Distance between TNS1 and TNS2*

k = *Distance power, when set equal to 0, the routing of traffic is not affected by distance*

Figure 7.23:

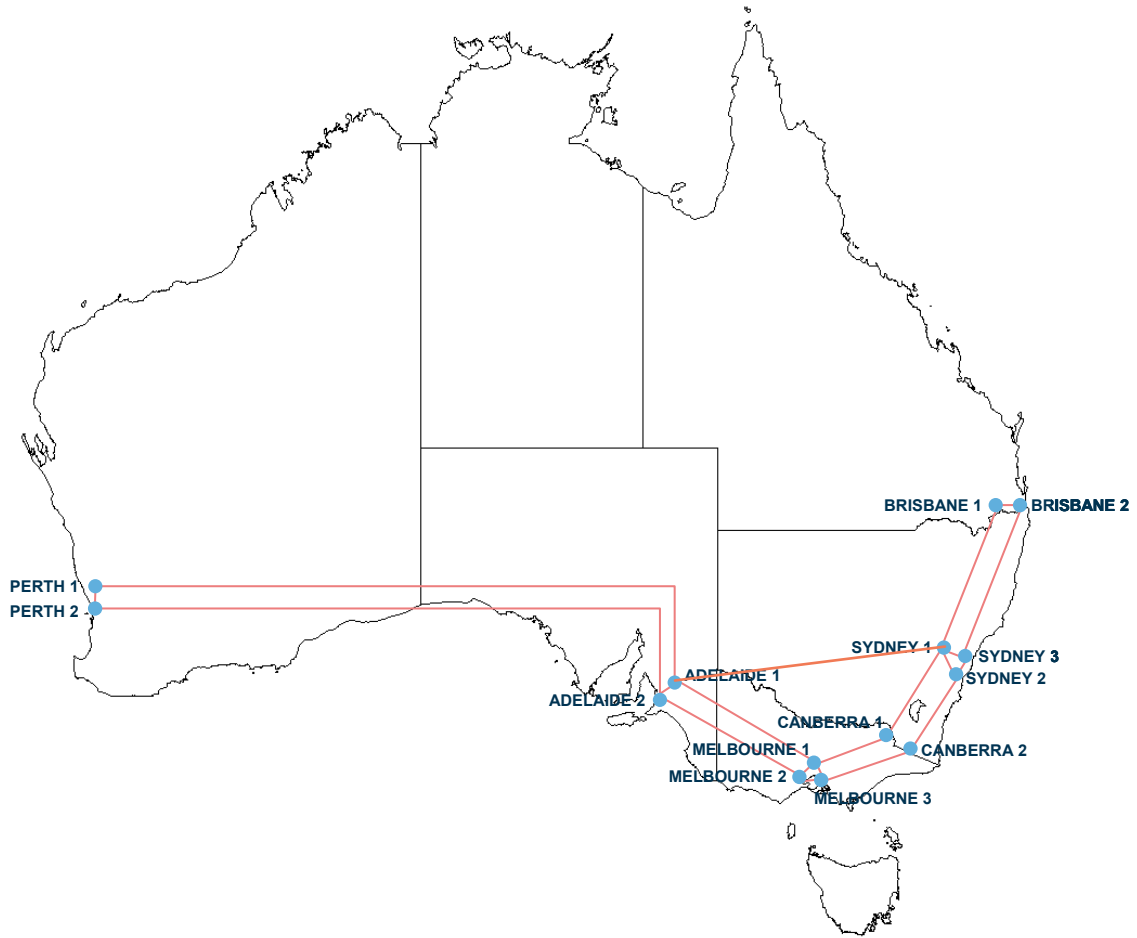
Formula for TNS
gravity model

[Source: Analysys]

The output is the percentage of traffic at each particular TNS that flows to each of the other TNSs.

7.9.4 TNS physical ring dimensioning

The TNSs are connected together in a series of physical rings, as outlined in Figure 7.24.



Note: the red lines indicate the route of the inter-state transit links. The Core module uses the actual road/railway distances when calculating the cost of the deployed trench, duct and fibre.

Figure 7.24: TNS rings modelled [Source: Analysys]




Figure 7.25: TNS rings modelled – physical routes mapped [Source: Analysys]

The physical rings are deployed along roads/railways, and their length takes into account the real route distance according to GIS mapping calculations.

These physical rings are dimensioned according to the inter-switch traffic calculated for each of the logical links that each part of the ring needs to carry. For example, the physical ring linking Adelaide to Melbourne also needs to carry the traffic that comes from Perth whose destination is Melbourne, Canberra, Sydney and Brisbane.

An ADM unit is required at each TNS for each ring which uses the TNS. The capacity of the ADM is dimensioned according to the required total capacity on the SDH ring to which it serves.

A fibre regenerator unit is dimensioned for every 80km of fibre sheath length. The core module calculates this requirement for each physical fibre link.

DWDM equipment

At the transit level, the Core module deploys DWDM equipment, which is used to increase bandwidth over existing fibre-optic backbones by transmitting data over different wavelengths in the same fibre.

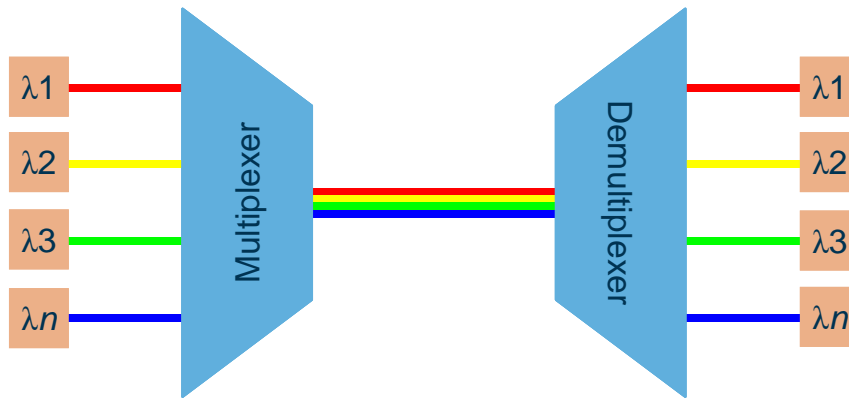


Figure 7.26:
DWDM equipment
[Source: Analysys]

The DWDM systems include multiplexing and de-multiplexing equipment. The former is required to take individually-encoded, multiple wavelengths and combine and couple them into a single optical fibre. At the receiving end of the fibre, the inverse process is performed using the de-multiplexer.

Whilst there is little incremental cost to deploying multiple fibres rather than a single fibre, DWDM presents a cost saving when regenerator equipment is taken into account. Regenerator equipment is required for each fibre.

7.10 Special island solutions

Incorporating special bespoke solutions required for certain island locations into the main network design algorithms would prove a costly exercise from a computational point of view. Consequently, these solutions are modelled explicitly on the 'NwDes.5.Islands' worksheet.

Special solutions are only modelled for those islands which have a LE located on them. If an island is part of a large ESA and is part of the sample, then the rural deployment algorithm in the geoanalysis and access network module can determine whether the clusters of demand on that island should be served by a wireline or wireless solution. Given the distance between the island and the remainder of the ESA on the mainland, such clusters would then be provisioned with either satellite or microwave backhaul to the LE of that ESA.

Special island solutions are determined at two distinct levels in the network – at the LE–PoC level (island LEs are removed from the possible list of PoCs), and at the LAS–TNS level (as two LAS exist on Tasmania).

With regard to the former of these, the type of network solution is user-defined in the Core module such that a microwave or a satellite solution is deployed. When an alternative wireless or satellite solution is deployed, the corresponding wireline distance as previously calculated in the Core module is subtracted to ensure that there is no double counting of transmission links.

7.10.1 Satellite solutions

The satellite option is deployed for large islands (namely Lord Howe Island, Christmas Island and Cocos (Keeling) Islands), which are located at a distance greater than the physical distance of a microwave link. A satellite solution requires the deployment of a satellite earth station in the core network.

7.10.2 Microwave solutions

The microwave solution requires the deployment of microwave towers at either end of the link, with microwave hop towers deployed where the link distance is greater than the maximum physical link distance of a microwave signal.

7.10.3 Submarine cable solutions

A submarine cable solution can be deployed if demand is sufficient – microwave and satellite solutions are not an appropriate backhaul solution for large demand volumes.

The Core module captures the existing 240km submarine cable deployment, between the mainland and Tasmania, with two separate cable paths across the Bass Strait.

These submarine cables are dimensioned to carry the traffic on the LAS ring incorporating the two LAS nodes on Tasmania.

A landing station is modelled to be deployed at either end of the link. The sub-component costs associated with a landing station, such as fibre termination equipment, optical fibre drivers and receivers, power supply, back-up power and air conditioning are included in the cost of a landing station.

7.11 Evaluation of access and core network overlap

This section outlines the process by which Analysys has determined the level of overlap between the modelled access and core network routes, i.e. the degree to which access and core network routes are deployed in the same trench.

From a cost perspective, it is important to determine the amount of incremental trench required in the network, as well as the percentage of the trench assets that are shared between both between the different levels in the core network and between the core and access networks.

It is not possible, from the LRIC Excel modelling, to directly determine the level of trench sharing between the different levels of the core network or between the core and access parts of the network. Consequently, a parameterised approach has been taken in the LRIC model.

To inform these parameters, Analysys has mapped the core links between each level in the network (LE, PoC, LAS, TNS) using the road and railway network in Australia. The MapInfo StreetPro dataset for roads and railways has been combined with a route-plotting tool called RouteFinder²⁹ in order to generate the shortest distances between each of the nodes using actual routes. Using this method, the LE–POC minimum spanning trees, PoC rings, LAS rings and TNS rings were plotted.

The default trench sharing values in the LRIC model are based on this analysis.

7.11.1 System requirements

GIS tools have been used to process the data and to plot the street-based network deployment routes. The following tools and datasets have been used to plot the network routes and perform the overlap analysis:

- MapInfo 7.5 and above
- RouteFinder 3.41 For MapInfo
- MapInfo StreetPro Australia 2007
- MapBasic (version 4.5 and above) – required for the preparation of the street network.

The remainder of this section discusses the following:

- An outline of the methodology used to determine the overlap between the core and access networks
 - preparation of the street network
 - mapping of the core network routes
 - definition of the access network
 - overlap analysis
- the results of the overlap analysis and the sensitivity analysis performed
- the caveats placed upon the analysis.

The processing required within MapInfo is described in more detail in section 7 of the *Instructions for key processes in geoanalysis for the fixed LRIC model*.

7.11.2 Methodology: preparation of the street network

To map the physical trench routes required for the core network, Analysys has used the MapInfo StreetPro Australia network map together with the route plotting software, RouteFinder. In order to ensure that the StreetPro Australia map files were in an appropriate format usable by RouteFinder, two stages of pre-processing was required:

²⁹ The standard version of RouteFinder was used (<http://www.routeware.dk/routefinder/routefinder.php>).

- Firstly, the StreetPro dataset was filtered in order to reduce the number of routes under consideration. This was required to enable reasonable computation times, as well as taking consideration of the data limit of RouteFinder Standard (500 000 entries)
- Secondly, the StreetPro dataset was processed in order to convert it into a format usable by RouteFinder.

7.11.3 Methodology: mapping of the core network routes

Analysys has mapped the physical core network routes onto the actual Australian road and railway network for each of the core network levels based on the logical mapping of these routes contained within the Core module:

<i>Network level</i>	<i>Type of network topology</i>	<i>Location in Core module</i>
LE	Minimum spanning tree	In.Nodes
PoC	Ring structure	In.Nodes
LAS	Ring structure	NwDes.3.Reg.Nodes
TNS	Ring structure	NwDes.4.Core.Nodes

Note: the LE and PoC routes are based on the two bridging node TSP algorithm for up to the 5254 individual ESA-copper centres combinations (possible LE location).


Figure 7.27: Source of core network route mapping [Source: Analysys]

Additional trench from core nodes

When mapping the core nodes onto the StreetPro network, RouteFinder translates the nodes to the nearest road segment. This translation means where a node does not exactly sit on a route, a small distance may be missed for each node. This distance is calculated using the ‘Co-ordinate Extractor’ tool in MapInfo and accounted for in the overlap analysis.

Core network

The generated route map for each of the network levels is shown in the following diagram:



Note: dark blue lines are the LE–PoC routes; green lines are the PoC–ring routes; light blue lines are the LAS–ring routes; pink lines are the TNS–ring routes.

Figure 7.28: Core network routes [Source: Analysys]


7.11.4 Methodology: definition of the access network

In the geographical analysis supporting the access network deployment, only a sample set of 200 exchange service areas (ESAs) in Australia were modelled. This overlap analysis requires that consideration be taken of the access network within every ESA. Consequently, Analysys has defined a proxy access network area using two measures:

- The built-up area (BUA) layer shape from the StreetPro Australia dataset – this is MapInfo’s generalised representation of areas with a relatively high level of infrastructure
- A 4km-radius buffer area around each local exchange copper centre location – this is Analysys’s current assumption on what would be a reasonable distance to expect sufficient CAN trench deployment to support IEN network. It equates to a distance where copper easily

supports an access network. A sensitivity analysis has been performed on this parameter as it has a significant impact on the trench sharing percentage result

These areas form a shape against which the network routes for the different levels in the core network are assessed. The access network proxy shape is shown in the diagram below.



Regions of red identify the shape of the proxy access network

Figure 7.29: Access network proxy layout [Source: Analysys]

Analysys has performed a sensitivity analysis on the buffer size around each of the LEs as it has a significant effect on the overlap results. Results have been generated for 0km (i.e. just the BUA regions), 1km, 2km, 3km, 4km and 5km buffer radii. The square kilometre areas covered by these shapes are shown in the table below.

Shape	Area covered (km ²)
BUA	14 072
BUA + 1km buffer around LEs	29 423
BUA + 2km buffer around LEs	75 476
BUA + 3km buffer around LEs	152 232
BUA + 4km buffer around LEs	220 685
BUA + 5km buffer around LEs	397 848

Figure 7.30: Proxy access network shape areas [Source: Analysys]

7.11.5 Methodology: overlap analysis

IEN overlap

The overlap of the IEN routes allows the determination of how much core distance is incremental. There are two sets of incremental distances to consider:

- Trench (and duct) may be shared *intra* core network level, for example, parts of the LE–PoC pathways may be deployed along the same route, therefore could share trench:




Figure 7.31: Example of intra network level trench sharing [Source: Analysys]

- Trench (and duct) may be shared *inter* core network levels, for example, parts of the PoC-rings may be deployed along the same routes as LE–PoC routes:

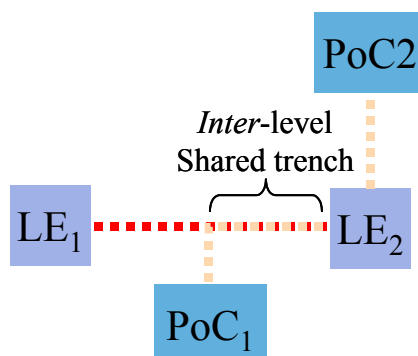


Figure 7.32: Example of inter network level trench sharing [Source: Analysys]

Consequently, in order to calculate the total incremental core trench required:

- The actual routes for each network level were combined using the ‘Combine’ MapInfo function, removing any trench double-counting associated with intra-level trench sharing

- The LE routes were buffered in MapInfo and any overlap with the PoC routes was removed and the distance of the incremental PoC trench could then be subsequently calculated
- The LE and incremental PoC routes were then used to similarly remove any overlap with the LAS routes
- The LE, incremental PoC and incremental LAS routes were then used to similarly remove any overlap with the TNS routes

We do not assume sharing of fibre on these routes as a fibre sheath needs to connect two nodes. The distances between nodes (i.e. not considering intra- or inter-level IEN sharing) is used to uplift fibre sheath distances from crow-flies to road-distances.

CAN-IEN overlap

In order to determine the overlap of the core network routes with the access network, each incremental core network level was compared against the proxy access network shape. The shape of the proxy access network is removed from the incremental core network route to determine the distance of the incremental core network outside of the access network.

The level of trench that is shared with the access network is calculated as the delta between the original incremental trench distance and the resultant incremental trench distance.

Consequently, the following sets of data are produced for each level in the core network:

- % trench shared with other levels in the core network (described above)
- % trench shared with access network.

The outputs are used in the core module to calculate the distance of trench required for the IEN outside of the CAN area. The distance of IEN within the CAN area is separately calculated to prevent double counting of trench distances. Instead this distance is used to allocate a proportion of trenching costs from the CAN to the IEN within the cost module.

Figure 7.33 shows the output of the core network routes overlapping the proxy access network (dark green) in the urban area of Melbourne. This map output shows several important effects:

- there is a very high level of overlap between the core and access networks in urban areas
- many of the buffered LEs overlap – this reduces the incremental effect of the buffering in urban areas
- the urban areas are dominated by PoC-ring topology
- the TNS routes are mainly deployed in pre-existing trench (i.e. trench laid at either the LE, PoC or LAS levels) – demonstrated by only small incremental pink routes plotted.




Figure 7.33:
Plotted routes in
urban Melbourne.
Area shown is
340km x 210km
[Source: Analysys]

Figure 7.34 shows the output of the routes overlapping the proxy access network in a rural area (in Queensland). This map output demonstrates the very small level of overlap of access and core networks in rural areas. However, it also shows that whilst the distances of the individual routes are far longer than in urban areas, the number of these routes is far smaller.




Figure 7.34:
Plotted routes in
rural Queensland
Area shown is
340km x 210km
[Source: Analysys]

7.11.6 Results of access and core network overlap

The results of the overlap analysis is presented in the table below.

<i>Network level</i>	<i>Total distance (km)</i>	<i>Incremental distance (taking into account intra-level trench sharing) (km)</i>	<i>Total incremental IEN distance (km)</i>	<i>Overlap of incremental IEN distance with access network (BUA + 4km LE buffer) (km)</i>	<i>Percentage incremental to access network</i>
LE level	136 884	114 464	114 464	38 145	33%
PoC level	54 799	43 125	22 662	9 677	43%
LAS level	37 212	25 587	10 708	3 161	30%
TNS level	13 337	12 552	2 635	314	12%
Total	242 232	195 728	150 468	51 296	34%

Figure 7.35: Overlap analysis: core network overlap with access network results [Source: Analysys]

A sensitivity analysis of the LE buffer level when defining the proxy access network is presented below.

Sensitivity analysis

The sensitivity of the results to the buffer size around each of the LEs when defining the proxy access network shape is presented below:

<i>Network level</i>	<i>BUA plus buffered LE</i>					
	<i>BUA</i>	<i>1km buffer</i>	<i>Network level</i>	<i>BUA</i>	<i>1km buffer</i>	<i>Network level</i>
LE level	4%	10%	LE level	4%	10%	LE level
PoC level	21%	23%	PoC level	21%	23%	PoC level
LAS level	12%	15%	LAS level	12%	15%	LAS level
TNS level	4%	6%	TNS level	4%	6%	TNS level
Total	7%	12%	Total	7%	12%	Total

Figure 7.36: Percentage of IEN incremental to access network [Source: Analysys]

The 4km buffer results have been selected as the current assumption on what would be a reasonable distance to expect sufficient CAN trench deployment to support IEN network. These results demonstrate a 34% (99 172km) incremental core network beyond the CAN area.

Caveats

A number of caveats should be raised around the access and core overlap analysis:

- There is potential for more inter-core trench network sharing: trench deployed at a different network level does not take into consideration trench deployed at previous levels, RouteFinder simply calculates the shortest route between two points.
 - increased inter-level sharing in the core network will serve to increase the percentage shared with access number, as it will lower the total incremental trench deployed in the core network.
- An uplift multiplier of 1.7 has been used to estimate road distance versus straight-line distance for the additional trench deployments (for when RouteFinder moves nodes onto the nearest road segment), based on the results of the LE-POC route analysis.
- The multi-ring travelling salesman algorithm has been used to calculate the efficient ring structures for the PoC rings. The algorithm assumed straight-line distances between the nodes. This route analysis maps the PoC rings onto actual roads, resulting, in some cases, in spur topologies rather than ring structures being generated.

7.11.7 Implementation of access and core network overlap

To capture the cost saving arising from the use of CAN trenches by duct used for the IEN, we calculate utilise the overlap analysis results in the Core module. The modelling approach adopted is based on deploying IEN duct in existing CAN trench and allocating a proportion of the CAN trench cost to the IEN.

In the core module, the distance of duct deployed within and outside of CAN areas is explicitly calculated, with both elements being costed. The cost of the underlying trench is calculated for volumes outside of CAN areas. Costs for trench within CAN areas are not calculated, instead an allocation of the cost of the CAN trench utilised is allocated to the IEN.

The proportion of cost allocated from CAN to IEN is a user input and explained in section 6.13 of the user guide. The default input is 50%.

The CAN buffer area is a user input, and can be varied. The default input is 4km.

8 Core network design – next-generation network

This section describes the next-generation network design of, and algorithms implemented to construct the core network. In implementing both the modern and the next-generation network design, algorithms developed can accommodate both approaches³⁰. Consequently, the worksheet naming terminology takes a generic approach for the two technologies.

In the NGN, traffic is transported as IP-based packets in the core network. Voice traffic is converted to IP at MSANs that are connected to customers, and at the media gateways (MG) that interfaces with traditional TDM switches, under the control of a soft switch. They connect, in turn, to the networks of other operators.

A high-level structure of the core NGN architecture is shown in Figure 8.1 below.

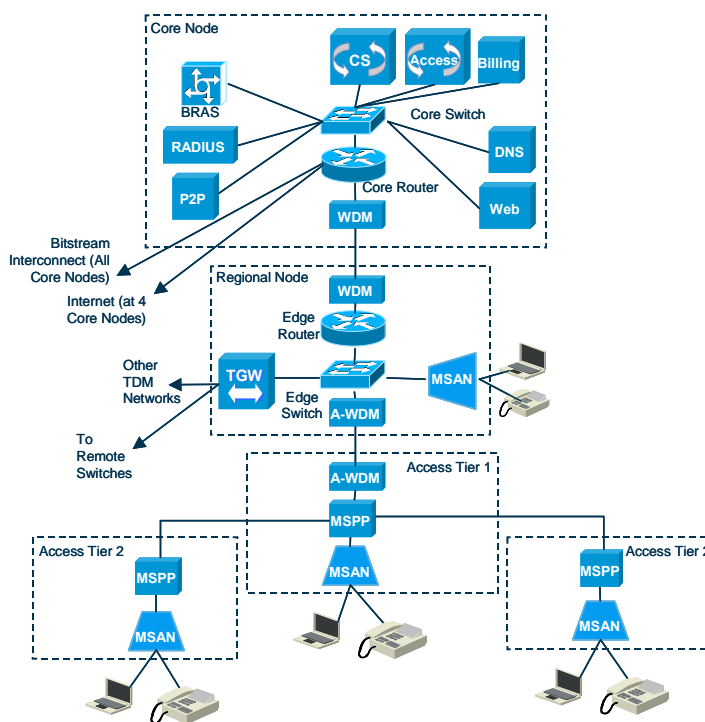


Figure 8.1:
High-level overview of the
core NGN [Source:
Analysys]

Each of the core nodes (TNS/MTH equivalent) is connected to a number of regional nodes (LAS/LTH equivalent). These regional nodes, in turn, connect to the access nodes arranged in two tiers, with access tier-two sites (pillar equivalent) being parented by an access tier-one site (LE equivalent). Within each access node, an MSAN is used to connect the copper pairs for each

³⁰ An NGN network design is implemented by setting a number of the geotypes in the 'MSANs deployed in geotype' parameters to 1. Defined in the named range *Geotype.served.by.MSAN.equipment* named range on the 'Scenario' sheet of the Cost module, passed to the 'In.Control' sheet of the Core model.

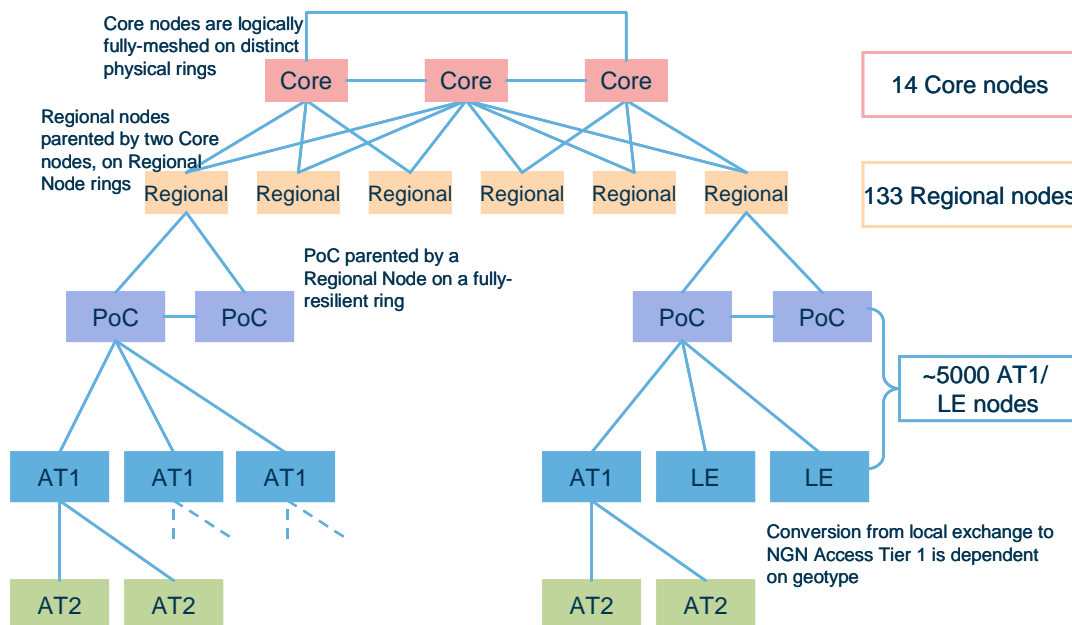
customer, with the MSAN then converting the voice and DSL into a single IP-based backhaul to the regional and core nodes. A combination of ring topologies and resilient links is used for interconnecting the different nodes in the network. TDM-based voice traffic from other networks is delivered to the Trunk Gateway (TGW), which translates the TDM-based voice to IP for transit over the NG core network. Traffic from the existing remote switching units (LEs), which are not included as part of the NGN, is also connected to the TGWs.

At each of the regional and core nodes, there are Ethernet edge switches to aggregate the traffic and Ethernet routers for routing between regional and core nodes. Within a core node, there is a soft switch to oversee the voice traffic and also equipment specific to DSL services (i.e. BRAS, mail servers, user authentication (RADIUS) servers and DNS). Billing equipment is also included at the core nodes.

DSL traffic is carried on the same network as voice traffic, as DSL traffic going to, or coming from, Internet peering locations, or directly to Internet service providers (ISPs) as part of the bitstream service.

The network also includes an underlying transport architecture consisting of a mixture of Ethernet over SDH – supported by multi-service provisioning platforms (MSPP) – and wave division multiplexing (WDM) systems.

As depicted in Figure 8.2, the topology of the nodes in the NG core network follows that of the modern network deployment, with the addition of the Access Tier 2 nodes – which are located at the pillar location in metropolitan areas.



Note: AT2: Access Tier 2; AT1: Access Tier 1; LE: Local Exchange; PoC: Point of confluence; Regional: Regional nodes; Core: Core nodes.

Note: a PoC is an AT1 on a resilient ring.

Figure 8.2: Network topology [Source: Analysys]

The network design follows NGN design principles, but accommodates the geographical characteristics of the existing network. A network consisting of four layers – core, regional, Access Tier 1 and Access Tier 2 – has been designed.

The number of nodes in the NGN topology currently aligns with the modern network topology – Access Tier 1 nodes are located at the LE locations, regional nodes are located at the LAS locations, core nodes are located at the TNS locations. The Core module does contain the flexibility to enable different routing of traffic between nodes, and hence the number of nodes in the NGN could be rationalised relative to the current modern network modelled.

The underlying transport architecture consists of a mixture of Ethernet over SDH MSPP elements and WDM systems.

As a base case for NGN deployment, the LRIC model could assume that NGN design is only implemented in urban areas. In rural areas, like the modern network design, the TDM design algorithms are implemented up to the LAS layer, at which traffic is inter-switched at a TGW onto the IP core.

8.1 Demand drivers

The demand loading on the modern network is calculated on the ‘Dem.Calc’ worksheet.

The main demand drivers used to dimension the various network elements are derived from the total service volumes that are output from the Demand module (refer to section 3).

The demand calculation for the NGN traffic follows the same calculation flow as for the modern network design, with the exception that the output of the demand is in terms of a busy-hour kbit/s load. In an IP network, voice requires a bit rate of 135kbit/s (this is in fact greater than under a TDM architecture, in which voice is transmitted at 64kbit/s); this is based on G.711 10ms packetisation over Ethernet, in which additional bandwidth is required for IP overheads.

It is worth noting that under NGN deployment scenarios, the provisioned backhaul per NGA subscriber has an important impact on the dimensioning of the core network. See sections 5.9 and A.6 of the user guide (‘In.Network’ worksheet). This provisioning differs from the modern network approach where backhaul is provisioned on a per-rack basis.

8.2 Routing factors *(Sheet DemCalc, Rows 273-375)*

A separate set of routing factors is required for the NGN traffic, however, these routing factors are based upon the same traffic routes as utilised by the modern network traffic.

8.3 NGN transmission

The NGN model considers Ethernet traffic in terms of Mbit/s. It dimensions four transmission speeds – 10Mbit/s, 100Mbit/s, 1GE and 10GE. By using a similar formula to the one used in the transmission calculations for the modern network, the NGN transmission algorithms consider the number of assets of a single speed required to meet the bandwidth requirement. This calculation also takes into account the cost threshold associated with deploying these transmission assets. It is estimated that up to four lower-speed units can be deployed before it is more cost efficient to deploy a higher-speed link.

8.4 Access Tier 2 (*Sheet NwDes.1.Access*)

In urban areas, it is assumed that MSAN equipment is deployed at the location of the pillar to allow for VDSL2+ services to be provided over the much shorter link-lengths. This pillar location is the same as for the modern network design. Consequently, data regarding the number of pillars and distance of links from pillar to LE (equivalent to the Access Tier 2 to Access Tier 1 routes) is linked from the CAN module.

Unlike the calculations that are performed at each of the other node levels, the Access Tier 2 calculations are performed at an aggregate level (at the Access Tier 1 level of detail).

8.4.1 Equipment dimensioning (*Sheet NwDes.1.Access Row 5271-10529*)

The number of MSANs is calculated for each of the Access Tier 2 locations (calculated in an aggregate form at the Access Tier 1 level of detail).

The MSAN equipment modelled is assumed to be able to handle multi-platform traffic, i.e. PSTN, ISDN and xDSL traffic. The MSAN equipment is dimensioned in terms of ports, shelves and racks, and is modelled to reflect the Alcatel ISAM 7302 equipment that is currently being deployed in its NGN network.

Each shelf can handle 18 line cards. It is assumed that a particular line card can only handle similar traffic, i.e., the number of PSTN, ISDN and xDSL line cards is dimensioned separately, but can be plugged into the same shelf. A PSTN line card can take 48 PSTN ports; an ISDN line card can take 24 ISDN-BR ports, or 16 ISDN-PR ports; and an xDSL line card can take 48 xDSL ports.

An MSAN unit can have four shelves per rack.

8.4.2 Transmission dimensioning (*Sheet NwDes.1.Access Row 10533-15791*)

The MSAN equipment has Ethernet interfaces for backhaul purposes. The Access Tier 2 nodes are connected to a parent Access Tier 1 nodes by means of a fibre ring, the speed of which is dimensioned by the total demand on the fibre ring.

AT2 equipment is assumed to be deployed at the pillar location. Each AT2 is linked to a parent AT1 node (which is located at the LE location in the modern network design). The length of the AT2-AT1 transmission links is assumed to be equal to the distance of the pillar-LE links as determined from the access modelling for the modern network.

8.5 Access Tier 1 (*Sheet NwDes.1.Access*)

Each Access Tier 2 node is connected to an Access Tier 1 node via fibre rings. The Access Tier 1 nodes host subscribers and also act as aggregation points for Access Tier 2 nodes for connectivity to the regional and core nodes.

8.5.1 Transmission dimensioning (*Sheet NwDes.1.Access Row 10533-15791*)

The multi-service provisioning platform (MSPP) at the Access Tier 1 node is an aggregation device for all MSAN Ethernet links hosted by the Access Tier 1 node.

Access Tier 1 nodes, not already designated as PoCs, are connected to PoC nodes using direct Ethernet links.

8.6 Aggregation nodes (PoCs) (*Sheet NwDes.2.PoC*)

Similar to the modern network design, Access Tier 1 nodes are aggregated at PoCs that are connected to a parent regional node (which are located at the LAS location).

When moving to an NGN architecture, the LRIC model accounts for the fact that some parts of the network will still be using TDM technology. Consequently, at the PoC level – at which a mixture of technologies is backhauled – the LRIC model deploys one of three possible choices of backhaul technology:

- TDM over SDH – used only where each PoC on the PoC-ring utilises TDM technology
- Ethernet – used only where each PoC on the PoC-ring utilises Ethernet technology
- Ethernet over SDH – used where the PoCs on the ring share a mixture of TDM and Ethernet technology. SDH boxes are assumed to have cards with Ethernet interfaces, particularly for 10/100 Ethernet.³¹

Regarding Ethernet over SDH, the generic framing procedure (GFP)³² allows Ethernet to be mapped on to different VCs: n×VC-12 (2Mbit/s), n×VC-3 (48Mbit/s) or VC-4 (149Mbit/s). In calculating the required number of VCs, the required bandwidth should be considered rather than

³¹ This implemented approach follows the BT approach taken in the UK, and is deemed to be realistically applicable to Australia.

³² Ethernet frames sent on the SDH link are passed through an "encapsulation" block (typically GFP) to create a synchronous stream of data from the asynchronous Ethernet packets.

the bearer capacity of the interface, i.e., if only 30Mbit/s of a 100Mbit/s Ethernet interface is required. The capacity of the SDH transmission levels (STM-x) in terms of VCs is shown in Figure 8.3 below:

<i>SDH transmission level</i>	<i>VC-12 capacity</i>	<i>VC-3 capacity</i>	<i>VC-4 capacity</i>
STM-1	63	3	1
STM-4	252	12	4
STM-16	1008	48	16
STM-64	4032	192	64

Figure 8.3: VC capacity of STM-x SDH interfaces
[Source: Analysys]

The Core module calculates the number of VC-3s needed to handle the required bandwidth. This information is used to calculate the appropriate STM-x bearer necessary to handle the number of VC-3s.

8.7 Regional nodes

The regional nodes are assumed to be located at the site of the local transmission hub in the modern network design. In addition to providing the TGW interface to the TDM network, the regional node provides connectivity to two of the 14 core nodes via rings. The links carry DSL traffic and voice connectivity to other TGWs and MSAN-hosted subscribers elsewhere in the network. These regional node rings are of the same topology as the LAS rings (the topology of the LAS rings are outlined in section 7.8.2).

The Core module assumes that the regional nodes route all traffic back into the core, i.e. there is no local routing of traffic.

8.7.1 Equipment dimensioning

The regional nodes contain two major pieces of equipment, namely edge switches and edge routers:

- **Edge switches** (*Sheet NwDes.3.Reg.Nodes Rows 434-573*) – these elements aggregate traffic from the TGW and MSANs for delivery to and from the core and access nodes. They require 48 port cards to link to the MSANs and 12 port cards to link TGWs and edge routers. The capacity of a chassis is five slots for connectivity cards.
- **Edge routers** (*Sheet NwDes.3.Reg.Nodes Rows 576-715*) – these elements route traffic from the regional nodes to the core nodes using 1Gbit/s (two port) cards. The chassis unit has a capacity of 12 card slots and we assume that each edge router has a minimum of two of these cards.

Each router or switch requires chassis elements and port cards:

- **Chassis** – the capacity of a chassis is based on the number of card slots that it contains. Enough chassis are bought to house all the required port cards. The cost of a chassis includes all fixed elements such as control cards.
- **Port cards** – port cards contain different numbers of ports and support different data rates. The number of port cards of each type required at each router or switch is calculated by dividing the number of ports on a card by the number of ports required. We also account for minimum deployment at each node.

8.7.2 Transmission dimensioning (*Sheet NwDes.4.Core.Nodes*)

At the regional node location, all traffic is carried on the IP network; consequently, the Core module deploys the regional node rings using Ethernet transmission interfaces.

8.8 Core nodes

Core nodes are located in the same node locations as the TNS in the modern network. However, unlike the TNS within the modern network design which only handled voice traffic, the core nodes for the core NG network route both voice and data traffic.

8.8.1 Equipment dimensioning

The core nodes contain two major pieces of equipment, namely core routers and core switches:

- **Core routers** (*Sheet NwDes.4.Core.Nodes Rows 111-141*) – these elements route traffic around the core network ring and link to the regional nodes using 1Gbit/s (four port) cards. The chassis unit has a capacity of 15 card slots and we assume that each core router has a minimum of two of these cards.
- **Core switches** (*Sheet NwDes.4.Core.Nodes Rows 144-165*) – these elements aggregate traffic for delivery to and from the core routers, DSL data-related elements and softswitch elements. Lower-capacity 48-port electrical GE cards link to the softswitch call server and access gateway plus the DNS and RADIUS servers, as only signalling traffic is carried across these links. The BRAS, Web server and core router are connected via the higher-capacity 12-port optical GE cards. The capacity of a chassis is five slots for connectivity cards.

Softswitch (Sheet NwDes.4.Core.Nodes Rows 495-517)

The softswitch is also located at the core node. In addition to the chassis and port cards, the soft switch requires the following elements:

- **Trunk gateway (gateway to TDM switches)** – the number required is calculated as the total number of provisioned channels for busy-hour voice traffic from IP to TDM (and vice versa) divided by the number of voice channels on an STM-1 (1890) plus an allowance for uneven distribution of traffic. The trunk gateway is dimensioned in units of 1 STM-1.
- **Call processor (signal processing)** – the number required is calculated from BHCA, with each processor being capable of processing one million BHCA. There is a minimum of one call processor per main core node.
- **Access gateway (gateway to IP network)** – the number required is calculated as the total number of subscribers divided by the single access gateway capacity (100 000 subscribers).

Other network elements

A number of other core network elements are located at the core node:

- **DNS server** – one is required per core node
- **RADIUS (and accounting system)** – one is required per link between a BRAS and a core switch
- **BRAS** – the number required is calculated as the total number of DSL subscribers logged on at any one time divided by the capacity of a BRAS.

Where only one unit of equipment is deployed at a core node, it is assumed that in the event of failure, equipment at other nodes will provide fall-back.

8.8.2 Transmission dimensioning

The NGN core switches are assumed to be located in the same locations as the TNS deployed in the modern network. Consequently, they are assumed to follow the same physical ring structure as deployed for the modern network.

DWDM equipment is assumed to be deployed to carry the volumes of traffic between the ring segments.

This DWDM is modelled to be either metro, long-haul, extended long-haul or ultra long-haul systems depending on the distance over which the links are to be carried. Fibre regenerators are assumed to be deployed every 80 kilometres.

9 Costing calculations

The Cost module determines the capital cost of building the access and core networks, and the operating and maintenance cost of operating and maintaining it. The module annualises the capital expenditure using a tilted annuity calculation – the results of which, with operations and maintenance costs, are used in the determination of the LRIC cost for the declared services modelled. It furthermore includes parameters for sharing of trench and duct assets between the core and access parts of the network.

The Cost module allocates the network costs between the various lines that use the access network and between the services that use the core network, based on service routing factors and service volumes.

The network cost of the fixed network can be examined in the Cost module.

This chapter discusses the following issues in further detail:

- scenario control worksheet
- unit costs
- asset lifetimes
- weighted average cost of capital
- tilted annuity calculation
- cost allocation between platforms
- treatment of common and indirect costs
- output results.

9.1 Scenario control worksheet

The scenario control worksheet within the Cost module is the primary control sheet for the entire LRIC model. The important parameters controlling the appropriate year to be modelled and the level of NGN deployment (on a geotype basis) are located on this worksheet. It should be ensured that these parameters are only altered on this worksheet and not in other parts of the LRIC model. When changes are made to these parameters, it needs to be ensured that the other active modules (Core and CAN) are open so that a full recalculation can be performed.

Other inputs are in discussed in more detail in section 6 of the user guide.

9.2 Unit costs

The costs included as capital costs are typically those related to the acquisition of durable equipment or assets expected to be operative for several years in the network. Operating and maintenance (O&M) costs generally are recurring costs.

Capital and O&M costs directly related to the size of the network are considered incremental. Thus, if the network expands due to an increase in demand, new equipment would need to be purchased and installed leading to an increase in both capital and O&M costs.

The Cost module assumes that all assets are deployed at the beginning of the year to meet a required volume of demand.

Within the cost module, input price and trends are defined in nominal terms.

9.2.1 Equipment costs

The capital costs for a network element reflects the list price to purchase and install one unit of equipment. These inputs have been estimated from three main sources:

- information provided by the ACCC
- benchmark data
- Analysys's estimates.

The Cost module has currently been populated with domestic and international benchmark asset figures, with Australian-specific data included where available. Whilst the LRIC model produces a result using these benchmark costs, it can be easily updated should more appropriate Australian-specific cost data become available.

The Cost module includes a schedule of capital costs for each network element in nominal AUD dollars for 2007, and a price trend (from 2007 onwards) which enables the price of modern equivalent assets over time to be determined. This price evolution is an important input into the tilted annuity calculation, as explained in section 9.5. The unit cost inputs are located in the Cost module, in the 'UnitCost.Core' and 'UnitCost.Access' worksheets.

It should be noted that:

- The fibre costs implemented in the Cost module include jointing costs for long distance cables
- The NMS cost includes the cost of centralized network equipment, and the routers used to manage and monitor the transmission equipment. The NMS is a common asset rather than an incremental asset, and is explained in section 9.7.

More detailed cost inputs are provide on the UnitCost.Core and UnitCost.Access worksheets for the following assets:

- Building costs (for LE, AT1, LAS and TNS sites) can be varied by geotype
- Duct deployment costs vary by duct size, and can be separately defined for urban (own trench provided and for open trench) and for rural where ploughing deployment methods may be employed

- The duct asset separately identifies road-crossing volumes and ‘serving-pit to property boundary’ volumes allowing different cost inputs for these specific assets
- Copper cable cost vary by the number of pairs in the bundle, and are separately defined for 0.4mm, 0.64mm and 0.9mm gauge cable.
 - For the cable assets, the lead-in volumes are separately identified from the distribution volumes, so that assets recovered in product connection charges are can be removed from annual rental charges
- Copper jointing costs are built up from three components; jointing rate per pair and the cost of the enclosure.

Spares uplift

An uplift can be applied to increase capital cost for a network element to assume the purchase of additional spares within the defined lifetime of the asset.

The default value for the spares uplift is 0 per cent.

Installation uplift

It is assumed that operating costs associated with the network deployment of the asset are capitalised - an uplift to the unit capital cost of network elements is explicitly defined to accommodate this. Such costs include installation labour, tools, testing, etc.

However, some asset cost benchmarks implicitly include capitalised deployment costs (e.g. for trench, duct, copper cable and fibre cable).

The default value for the installation uplift is 15 percent.

9.2.2 Operating expenditure (Sheet UnitCost.Core Rows 432-634)

The operating and maintenance costs associated with certain network facilities are modelled as a percentage of capital cost per unit. This operating and maintenance input is specified according to the 2007 capital cost values, in subsequent years O&M costs per unit is calculated according to the operating and maintenance price trend.

The default O&M percentages of the capital cost range from 3 to 15 per cent, dependent on the asset type. Alternative O&M costs can be input by the user if required, which may be appropriate for NGN core assets, as well as for access assets on the UnitCost.Access sheet.

Default inputs are derived from an examination of operator accounts. Therefore the default O&M inputs employed are in effect an average over the lifetime in a network that was built over an extended period and has all ages of equipment present.

We recognise that certain O&M charges can be calculated on a bottom-up basis, for example energy costs of a platform based on the equipment's specification. However, we have not been able to model all O&M costs on a bottom-up basis and therefore believe it is more consistent to employ a top-down approach for all inputs. These O&M parameters are easily changeable, and therefore a bottom-up estimate can be applied to the model by the user or adjustments can be applied to the existing inputs.

9.3 Asset lives

The Cost module uses the following default values for the network asset economic lives:

Type of asset	Years
Building	50
Network equipment	10
Line card	5
Core Network equipment	20
Building equipment	15
Copper	25
Fibre	25
Trench	40
Duct	40
Business overheads	10
Licence	5
IT system	3

Figure 9.1: Current asset economic lifetime assumptions [Source: Analysys]

9.4 Weighted average cost of capital

The Cost module has a separate 'WACC' worksheet for adjustment to the WACC parameters.

The default parameter values are based on the ACCC's *Unconditioned Local Loop Service Pricing Principles and Indicative Prices – 2007-08 WACC parameters*, June 2008, with an adjustment to the risk free rate to take account of changing economic conditions. These parameters however, do not necessarily reflect the ACCC's current views on these parameter values.

9.5 Tilted annuity calculation

The Cost module annualises the capital expenditure using a tilted annuity calculation. This tilted annuity formula has been applied on the 'TA.Access' and 'TA.Core' worksheets.

The annuity method calculates an annualised charge for capital cost recovery, given the cost of capital of the entity, in such a way that the total (the cumulative annuity) is predictable over time. A flat annuity calculates a constant annualised cost per annum that, after discounting, fully recovers the investment and return on capital employed. Operating expenditures are added to the cost recovery in each year they are incurred.

The tilted annuity methodology has been implemented in order to reflect that the modern equivalent asset (MEA) price for many assets changes over time. In a tilted annuity, the annualised cost of recovering the investment and return on capital is tilted with the forecast price trend of the asset, subject to still fully recovering investment and capital employed. As such, tilted annuities are sometimes used as a proxy for economic depreciation, particularly where the output of the asset does not change significantly over the period. This is the case in fixed networks, more so than in wireless networks.

The annuity charge formula is calculated as follows:

$$\text{AnnuityCharge} = \frac{WACC - (\text{MEAPriceChange} + \text{TiltAdjustment})}{1 - \left(\frac{1 + (\text{MEAPriceChange} + \text{TiltAdjustment})}{1 + WACC} \right)^{\text{lifetime}}} \times GRC$$

In this formula, the lifetime applied may be the financial lifetime or the economic lifetime (if different).

By the nature of the implementation of a tilted annuity cost annualisation algorithm, only the capital cost in the current year needs to be modelled.

9.6 Cost allocation between platforms

Each cost is identified as being either:

- **incremental** – capital cost or operating and maintenance directly attributable to a service increment (access or core)
- **common** – shared across network platforms
- **indirect** – business overhead.

Certain costs associated with the access or core increments have to be re-allocated to the other increment, as explained below in section 9.6.1. Allocation of costs between platforms in the core increment are explained in section 9.6.2.

The treatment of common and indirect costs is explained in section 9.7.

9.6.1 Sharing of costs between the access and core networks

The Cost module explicitly accounts for asset sharing between the access and core networks in terms of:

- trench and duct assets shared between the access and core network; and
- exchange costs support the MDF, which is part of the access increment

Both sharing assumptions use a parameterised approach and result in a cost allocation.

The former sharing assumption is informed by geoanalysis which has compared the overlap of the plotted network routes with the extent of the urban areas in Australia – as outlined in section 7.11. These parameters are input in the core module as they are first used to calculate the appropriate length of routes which may overlap with the access network.

Certain core assets (such as the transmission links from the LPGS to the LE) are explicitly modelled in the Access network design algorithm. The cost associated with these assets is determined in the Cost module and passed on to the core network costing using an allocation matrix in the ‘Inputs.Access’ worksheet.

The resulting total cost per geotype is divided by the total demand volume by geotype to generate the per-unit network output.

9.6.2 Core cost allocation between platforms

As part of the core network design calculations, the percentage usage of certain assets (fibre, trench and duct) has been calculated at each of the transmission levels. These calculations enable the costs of these transmission links, and associated equipment to be directly attributed to the services (PSTN, ISDN, xDSL) that use them.

These calculations are performed on the ‘CostAlloc.Core’ worksheet.

This worksheet considers cost allocation at four separate levels:

- cost allocation between fibre and ‘Other duct services’ that use the trench and ducts deployed
 - this allows for the allocation of cost to duct rental
- cost allocation between SDH and ‘Other fibre services’
 - this enables cost to be allocated to other fibre rental services
- cost allocation between the service platforms, and ‘Other transmission services’ (such as leased lines)
- cost allocation between the PSTN, ISDN, xDSL and potentially³³ ‘Other platform services’.

³³ The CostAlloc.Core worksheet can be set up so that costs are allocated to ‘Other platform services’. In the default settings this is not used, with the exception of building costs from AT1, LE, LAS and TNS.

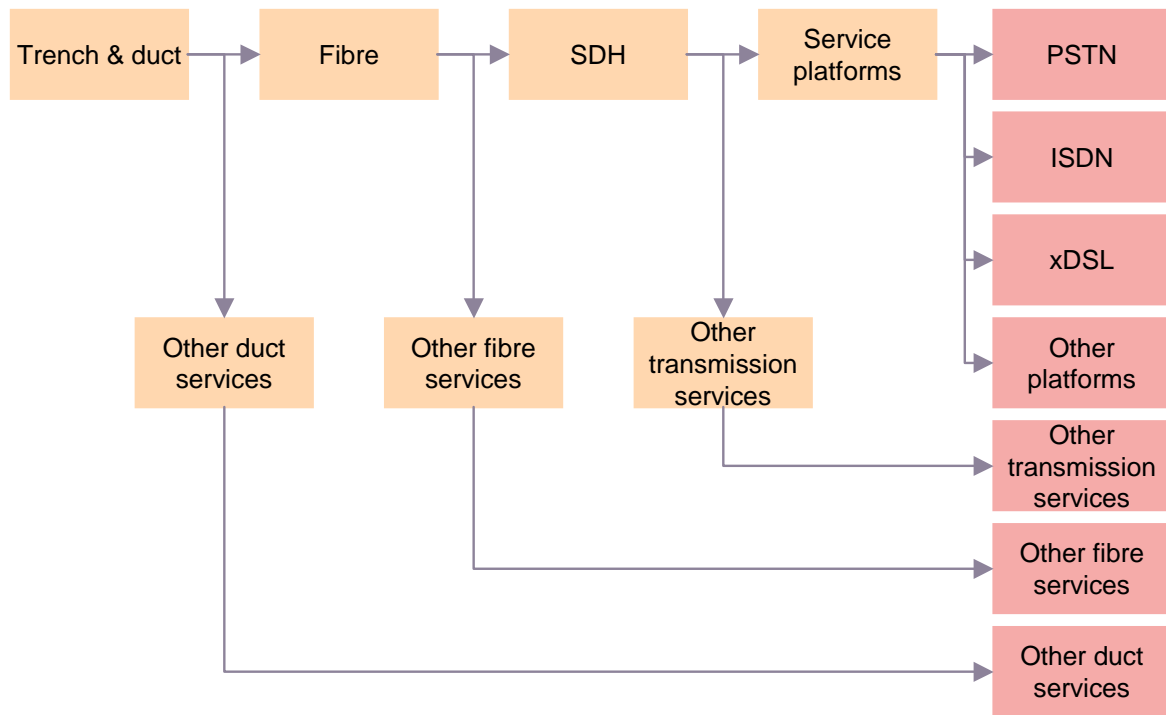


Figure 9.2: Cost allocation sequence [Source: Analysys]

In a similar manner, the use of exchange space (at the LE, LAS and TNS level) can be allocated to PSTN, ISDN, xDSL, and ‘Other platform services’ through the inputs on the I.Building.Core sheet. Other platform services is provided to capture costs of the exchange which support platforms not explicitly identified such as DDAS and mobile equipment, but does not include the SDH platform which is explicitly identified.

Capital and operating and maintenance costs directly attributable to the different platforms is used in the calculation of the LRIC result. The sum of the platform costs are divided by the total volume demand to generate the per unit network output.

9.7 Treatment of common and indirect costs

The common cost mark-up is performed in two separate stages. One mark-up is performed for costs common to core network elements only – identified with a ‘Core network equipment’ flag; and a second mark-up is for the costs common to all network elements – identified with an ‘All network elements’ flag.

The mark-ups are distributed across the platforms using an equi-proportional mark-up (EPMU) mechanism based on the level of incremental cost incurred by each platform.

The indirect (business overheads) costs are marked up on top of the combined total of incremental and shared costs of that network element.

Analysis of RKR reported figures are used to determine the value for business overheads. The methodology for determining the business overheads mark-ups are discussed in the following section.

9.8 Analysis of business overheads

The appropriate level of business overheads is an important input into the LRIC model. These business overheads represent organisational-level costs that are incurred by the business in order to generate business operations. They are appropriately marked-up on top of the determined network costs.

This section is set out as follows:

- Section 9.8.1 discusses Analysys’s treatment of Telstra’s RAF data.
- Sections 9.8.2 and 9.8.3 outline the operating cost and capital expenditure and associated business overhead results, including some aggregated data.

Specific reference has been made to RAF definitions presented in the ACCC paper “Telecommunications Industry Regulatory Accounting Framework (Record-keeping rules)”, issued May 2001.

9.8.1 Treatment of Telstra’s RAF data

Analysys has calculated the appropriate level of business overheads for the LRIC model using Telstra’s RAF CCA submission data from the Reporting Period June 2007.

Telstra submit data for their Retail business, External Wholesale business, and Internal Wholesale business separately. Consequently, retail costs and the associated retail business overheads are separately reported in Telstra’s RAF submissions, and so are excluded from the calculations of the appropriate level of business overheads uplift for the wholesale network assets. This division is shown in the figure below:

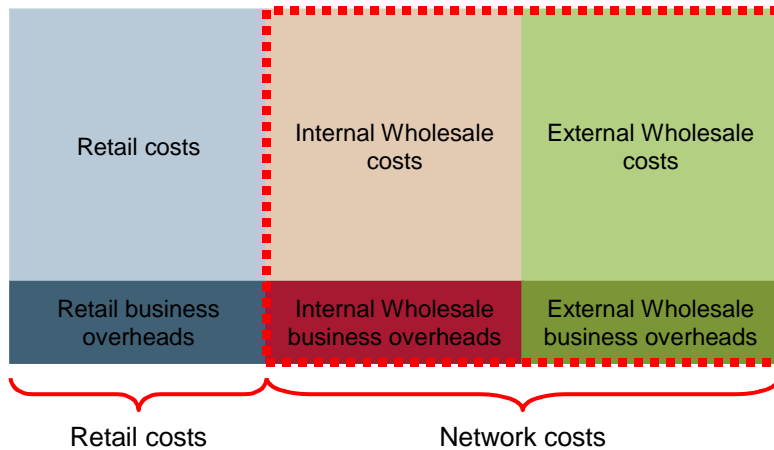


Figure 9.3: Division of costs [Source: Analysys]

Analysys has consequently utilised four specific sets of data from Telstra's RAF submission, in the determination of the appropriate level of business overheads:

- Capital adjusted profit (CAP) statement - external wholesale business.
- Capital adjusted Profit (CAP) statement - internal wholesale business.
- Fixed asset statement (FAS) - external wholesale business.
- Fixed asset statement (FAS) - internal wholesale business.

The CAP statements are used to determine the operating costs, and the FAS reports are used to determine the capital employed.

Exclusion of non-relevant wholesale services

The CAP statements and FAS reports are presented for the following wholesale services – listed separately for the internal wholesale business and for the external wholesale business:

<i>Internal wholesale business</i>	<i>Description</i>	<i>Included?</i>
Wholesale End User Access	Provision of access to the PSTN network not including retail specific costs	Y
Wholesale Local Calls	Non retail-specific costs for providing fixed network to fixed network calls within the carrier defined local call area	Y
Wholesale Domestic Long Distance	Non retail-specific costs for providing fixed network to fixed network domestic calls beyond the local call area, excluding international calls	Y
Wholesale International Long Distance	Non retail-specific costs for providing international calls originating either within or outside Australia	Y
Wholesale Domestic Leased Lines	Non retail-specific costs for providing all non-switched leased capacity solely within Australia including virtual private network and dedicated private circuits	Y
Wholesale Digital Data Carriage	Non retail-specific costs for providing access and carriage on a digital data network	Y
Wholesale ISDN Carriage	Non retail-specific costs for providing access and carriage voice and data on an ISDN network	Y
Wholesale Packet Switched Service	Non retail-specific costs for providing packet switched services designed for data communication between computers and data terminals which splits data into segments called packets	Y
Wholesale Fixed to Mobile	Non retail-specific costs for providing fixed network to mobile network calls terminating on domestically and internationally located mobile terminals	Y
Wholesale GSM Carriage	Non retail-specific costs for providing services, including access, originating on GSM networks	N
Wholesale CDMA Carriage	Non retail-specific costs for providing fixed services including access originating on CDMA networks	N
Wholesale Payphone Services	Non retail-specific costs for providing telephone calls from a carrier controlled pay telephone, whether owned, operated or leased. All other calls from non-carrier controlled payphones are captured under either local, domestic long distance, fixed to mobiles or International cost objects	Y
Wholesale Internet Services	Non retail-specific costs for providing Internet and related services including web hosting, web page design, SMTP and other e-mail services, electronic commerce and domain name administration services. Does not include local or domestic long distance calls for access to the Internet or carriage of Internet traffic	N
Wholesale Information Services	Non retail-specific costs for providing recorded message and other information sharing services, including community service calls, home shopping, chat services and electronic funds transfer, or as otherwise specified by the Commission	N
Wholesale Specialised Call Services	Non retail-specific costs for providing call services which involve subscriber customised functions in the handling, routing or changing of calls, including toll free, charge card, Freephone services, alternate and credit card-billing services, call waiting, call distribution and forwarding, call screening, UPT and virtual private network services, or as otherwise determined by the Commission	N
Wholesale Directory Services	Provision of a service to allow Freephone and Local rate customers number portability when changing carriers	N

Other Internal Wholesale Services	Non retail-specific costs for providing directory services including printed alphabetical or classified listings, electronic listing, through connect services, community, business, maritime, telex and facsimile directories	Y
<hr/>		
<i>External Wholesale Business</i>	<i>Description</i>	<i>Included?</i>
<hr/>		
Unconditional Local Loop (Declared)	A service is the use of unconditioned communications wire between the boundary of a telecommunications network at an end-user's premises and a point on a telecommunications network that is a potential point of interconnection located at or associated with a customer access module and located on the end-user side of the customer access module	Y
Domestic PSTN Origin/Term (Declared)	An access service for the carriage of telephone (ie. PSTN and PSTN equivalent such as voice from ISDN) calls (i.e. voice, data over the voice band) to a PoI from end customers' assigned numbers from the geographic numbers ranges of the Australian Numbering Plan and directly connected to the access provider's network	Y
Local Carriage Services (Declared)	A service for the carriage of telephone calls from customer equipment at an end-user's premises to separately located customer equipment of an end-user in the same standard zone	Y
GSM Origin/Term (Declared & Non-Declared)	An access service for the carriage of telephone calls (i.e. voice, data over the voice frequency band) to a PoI from end-customers to B parties, assigned numbers from the GSM number ranges of the Australian Numbering Plan and directly connected to the access provider's GSM network. This captures all wholesale provision of these services, both declared and non-declared	N
Digital Data Access (Declared)	An access service for the domestic carriage of data between a digital data Interconnect Terminal Point located at the access seeker's exchange or network facility and a network terminating unit or unimux or modem located at the customer's premises where the customer is directly connected to the access provider's network	Y
Transmission (Declared & Non-Declared)	A service for the carriage of certain communications from one transmission point to another transmission point via network interfaces at a designed rate on a permanent basis by means of guided and/or unguided electromagnetic energy. This captures all wholesale provision of this service, both declared and non-declared.	Y
Local Number Portability	Provision of number portability for retail PSTN customers wishing to maintain existing contact numbers while changing carriers	Y
Mobile Number Portability	Provision of number portability for retail mobile customers wishing to maintain existing contact numbers while changing carriers	N
Other External Wholesale Services	Other wholesale costs associated with providing services to external third parties	Y

Figure 9.4: Services covered in the internal and external wholesale reports (note: zero-cost services removed) [Source: Analysys, ACCC: Telecommunications Regulatory Accounting Framework (record-keeping rules, May 2001)]

The specific costs analysed are calculated by excluding the following services from the total costs, as they are not relevant to the services being modelled:

- Mobile-related costs (wholesale GSM carriage and wholesale CDMA carriage for the internal wholesale costs, and GSM origin/term and mobile number portability for the external wholesale costs)
- Wholesale Internet services
- Wholesale information services
- Wholesale specialised call services
- Wholesale directory services

Capital adjusted profit (CAP) statements – operating costs

The CAP statements are split into three sets of costs:

- Network costs (items listed as 4-3 in the CAP statements)
- Product and Customer costs (items listed as 4-2 in the CAP statements)
- Organisation costs (items listed as 4-1 in the CAP statements).

Network costs are considered for the wholesale services identified above. Depreciation is identified for network costs and is removed, as a return on capital is accommodated through the annuity calculation.

The Product and Customer costs are considered to be composed of a combination of business overheads and costs not relevant to the network services being analysed (termed ‘non-network’). The Product and Customer costs are comprised of the following costs: Installation, Marketing, Sales, Operator Services, Customer Support, Billing, Bad Debt Expenses, Interconnection Costs, International Settlement Costs, USO Payments, and Other Product Expenses. A description of these costs is stated in the table below:

<i>Line Item</i>	<i>Line code</i>	<i>RAF description</i>	<i>Relevance</i> ³⁴
Installation	4-2-01	Costs associated with the installation, facility removal or connection of equipment in relation to customer premise equipment. Excludes costs of major installation works at the main network	NON
Marketing	4-2-10	Costs arising from market and product research, promotions and advertising (both directly related to products and general) including mail outs	NON
Sales	4-2-15	Costs of direct and non-direct sales related activities including telemarketing, negotiations, technical information and support, account management	BO
Operator Services	4-2-20	Costs arising from the provision of services requiring operator assistance (for example directory services, operator connected calls, etc.)	NON
Customer Support	4-2-25	Costs of managing customer inquiries, complaints and disputes and credit management	BO
Billing	4-2-30	Invoice and receipt processing and dispatch	BO
Bad Debt Expenses	4-2-31	Debts written off or not expected to be collected. Debts written off but subsequently recovered are treated as Other Non-telecommunications revenue	BO
Interconnection Costs	4-2-40	Interconnection and usage expenses incurred to other national carriers for use of their networks	NON
International Settlement Costs	4-2-50	Out-payment expenses incurred to overseas carriers for call termination on their network	NON
Universal Service Obligation Payments	4-2-60	Payments to the government with respect to losses in net cost areas in accordance with Universal Service Obligation arrangements	NON
Other Product Expenses	4-2-70	Other product activities not already included such as printing of directories, materials management, product retirement costs, product-specific research and development etc	NON

Figure 9.5: Description of the Product and Customer costs [Source: ACCC: Telecommunications Regulatory Accounting Framework (record-keeping rules, May 2001)]

Analysys considers the Installation, Interconnection Costs, International Settlement Costs and USO Payments for wholesale services are likely to be recovered through separate charges, so have been excluded. Marketing, Operator Services, and Other Product Expenses are considered to not be relevant to costing of the network services modelled. Whilst these costs are considered non-network, they should be apportioned a percentage of the business overheads.

The Organisation costs are considered to be Business Overheads to the business. These Organisation Costs consist of the following costs: General Administration, Information Technology, Accommodation & Property, Other Non Communications Asset Costs, and Other Organisational Costs. A description of these costs is stated in the table below:

³⁴ Relevance categories; BO = Business overheads, NON = non-network, EXC = excluded.

<i>Line Item</i>	<i>Line code</i>	<i>RAF description</i>
General Administration	4-1-01	General corporate type costs including, corporate centre, legal, regulatory, accounting and finance, human resources and personnel, and non- product specific research and development
Information Technology	4-1-10	All non-network information system asset costs (including depreciation) and operational costs such as data entry, information technology support etc
Accommodation and Property	4-1-20	Costs (including depreciation) for non-network related accommodation and property including their maintenance, upgrade and the supply of basic services (i.e. power, water etc)
Other Non-Communications Asset Costs	4-1-30	Depreciation, amortisation and other costs related to non-communications assets not already captured including intangibles
Other Organisational Costs	4-1-40	All other organisational costs not already captured

Figure 9.6: Description of the Organisation costs [Source: ACCC: Telecommunications Regulatory Accounting Framework (record-keeping rules, May 2001)]

Consequently, Analysys has identified these costs as either network, non-network or as business overheads:

<i>Stated costs</i>	<i>Line code</i>	<i>Type of cost</i>
Network costs	4-3	Network
Product and Customer Costs	4-2	Mixture of business overheads and non-network costs
Organisation Costs	4-1	Business Overheads

Figure 9.7: Stated costs in the CAP statements [Source: Analysys & ACCC: Telecommunications Regulatory Accounting Framework (record-keeping rules, May 2001)]

The business overhead costs are allocated between the network and non-network costs using an equi-proportionate mark-up methodology:

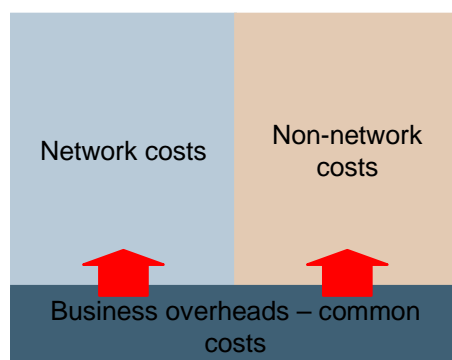


Figure 9.8: allocation of business overheads to network and non-network costs [Source: Analysys]

As stated previously, depreciation is separately reported for network costs, so it can be removed. For ‘Product and Customer’ and ‘Organisation’ costs depreciation is not clearly stated, so we have had to estimate adjustments, as explained below.

FAS reports– capital expenditures

The FAS are split into three sets of assets:

- Communications, Plant and Equipment (items listed as 2-1 in the FAS reports)
- Non-Communications Plant and Equipment (items listed as 2-2 in the FAS reports)
- Other Non-Current Assets (items listed as 2-3 in the FAS reports).

The Communications, Plant and Equipment costs is the network-related asset base.

The Non-Communications Plant and Equipment assets consist of Information Technology, buildings and improvements, and other assets. A description of these are provided in the table below:

<i>Line Item</i>	<i>Line code</i>	<i>RAF description</i>
Information Technology	2-2-01-4	All computer and related equipment not associated with the operation or management of the telecommunications network
Buildings and Improvements	2-2-03-4	All buildings including those under construction, owned by the company and the associated building alterations, extensions, structural attachments, improvements; and building plant and equipment such as electric light and power equipment, lifts and air conditioning which are considered an integral part of the building. Included are improvements on leased buildings
Other	2-2-05-4	All other plant and equipment not already included

Figure 9.9: Description of the Non-Communications Plant and Equipment assets [Source: ACCC: Telecommunications Regulatory Accounting Framework (record-keeping rules, May 2001)]

These are considered to be Business Overheads to the business.

The Other Non-Current Assets consist of long-term receivables, long term investment, intangibles, land and other assets. A description of these is provided in the table below:

<i>Line Item</i>	<i>Line code</i>	<i>RAF description</i>
Long term receivables	2-3-01-4	Balances owed by customers or other persons (for example, directors, officers, employees) not due within the next twelve months, net of any doubtful amount
Long term investment	2-3-05-4	Shares, options and other financial instruments intended to be held for more than twelve months
Intangibles	2-3-10-4	Identifiable assets having no physical existence, their value being limited by the rights and anticipative benefits that possession confers upon the owner, such as goodwill, trade marks, copyrights, patents, licenses and other assets of a similar nature
Land	2-3-20-4	All land (including freehold and perpetually renewable leasehold land) owned by the company. Includes land improvements but excludes buildings or other limited life structures
Other	2-3-25-4	Other assets not expected to be developed, utilised or received within the next twelve months and which have not already been included i.e. net long term inventories and provisions future income tax benefit and deferred research and development expenditure

Figure 9.10: Description of the Other Non-Current Assets [Source: ACCC: Telecommunications Regulatory Accounting Framework (record-keeping rules, May 2001)]

Of these items, only Land is considered to be an applicable business overheads. It has not been possible to associate the other categories as overheads to the network services being modelled.

<i>Stated costs</i>	<i>Line code</i>	<i>Type of cost</i>
Communications, Plant and Equipment Costs	2-1	Network
Non-communications, Plant and Equipment costs	2-2	Business Overheads
Other Non-Current Assets	2-3	Business Overheads (Land only: 2-3-20-4)

Figure 9.11: Stated costs in the FAS reports [Source: Analysys] [As above]

9.8.2 Operating costs

Using the stated method, the total network, non-network and business overheads costs are calculated for the fixed network. The adjustment to operating costs are shown in the table below:

	<i>Line code</i>	<i>Adjustments</i>
CAN Ducts & Pipes	4-3-01-2/3	0
CAN Copper cables	4-3-05-2/3	0
CAN Pair gain systems	4-3-15-2/3	0
CAN Other CAN	4-3-25-2/3	0
Local switching equipment	4-3-30-2/3	0
Trunk switching equipment	4-3-35-2/3	0
Other switching equipment	4-3-40-2/3	0
Inter-exchange cables	4-3-45-2/3	32,926,924
CAN Other cables	4-3-10-2/3	-32,926,924
Transmission equipment	4-3-50-2/3	0
CAN Radio bearer equipment	4-3-20-2/3	0
Radio bearer equipment	4-3-55-2/3	0
Data equipment	4-3-60-2/3	0
Mobile network and termination equipment	4-3-65-2/3	0
Customer equipment	4-3-70-2/3	0
Satellite equipment	4-3-75-2/3	0
International Network cables	4-3-80-2/3	0
International Network other systems	4-3-85-2/3	0
Other communications plant and equipment	4-3-90-2/3	0
Total		0

Figure 9.12: Network operating costs [Source: Analysys]

Product and Customer costs, and Organisation costs stated in the RAF submissions include depreciation. Consequently, adjustments should be made to the reported costs in order to remove the effects of depreciation. For Product and Customer costs, we do not find any noticeable level of depreciation therefore assume zero adjustment. For Organisational costs, we estimate a 10% deduction to Other organisational costs (4-1-40), which gives a weighted reduction of 0.29%. Non-network and business overheads costs, after adjustments are shown below:

<i>Item</i>	<i>Line code</i>	<i>Adjustment %</i>
Product and Customer – business overheads	4-2- 15; 4-2-25; 4-2-30; 4-2-31	0%
Product and Customer – non-network	4-2-01; 4-2-10; 4-2-20; 4-2-40; 4-2-50; 4-2-60; 4-2-70;	0%
Organisation costs (all business overheads)	4-1	-0.29%

Figure 9.13: Adjustments made to the Product and Customer costs and to the Organisation costs [Source: Analysys, Telstra 2007 RAF CCA submission - CAP statements]

The business overheads are distributed between the network and non-network components on an equi-proportionate basis, and the applicable uplift on the network operating costs for the business overheads is calculated as 59.77%

	<i>Total</i>
Business Overhead Uplift % ³⁵	59.77%

Figure 9.14: Business overheads uplift on opex
[Source: Analysys]

9.8.3 Capital expenditure costs

Using the stated method, the total network, non-network and business overheads assets are calculated for the fixed network. The adjustments to network capital employed are shown in the table below:

³⁵ Uplift % = Business Overheads attributed to network / Network costs

	<i>Line code</i>	<i>Adjustments</i>
Customer Access Network		
CAN Ducts & Pipes	2-1-01-1	
CAN Copper cables	2-1-05-1	
CAN Pair gain systems	2-1-15-1	
CAN Other CAN	2-1-25-1	
Switching Equipment		
Local switching equipment	2-1-30-1	
Trunk switching equipment	2-1-35-1	
Other switching equipment	2-1-40-1	
Inter-Exchange Investment		
Inter-exchange cables	2-1-45-1	2,466,025
CAN Other cables	2-1-10-1	-2,466,025
Transmission equipment	2-1-50-1	
Radio bearer equipment	2-1-55-1	
CAN Radio bearer equipment	2-1-20-1	
Other Systems and Equipment		
Data equipment	2-1-60-1	
Mobile network and termination equipment	2-1-65-1	
Customer equipment	2-1-70-1	
Satellite equipment	2-1-75-1	
International Network cables	2-1-80-1	
International Network Other	2-1-85-1	
Other Communications Plant & Equipment	2-1-90-1	
Total		0

Figure 9.15: Network capital employed [Source: Analysys]

Non-communications, Plant and Equipment, and Other Non-Current Assets costs stated in the RAF submissions are current cost Net Book Value numbers. Consequently, adjustments need to be made to these numbers in order to derive current cost Gross Book numbers. These adjustments are informed by data supplied by Telstra. These adjustments are identified in Figure 9.16.

<i>Item</i>	<i>Line code</i>	<i>Adjustment %</i>
Information Technology	2-2-01-4	170%
Buildings and Improvements	2-2-03-4	100%
Other	2-2-05-4	310%
Land	2-3-20-4	10%
Total business overhead assets		84%

Figure 9.16: Adjustments made to business overhead assets [Source: Analysys]

The business overheads are distributed between the network and non-network components (currently zero) on an equi-proportionate basis, and the applicable uplift on the network capital expenditure costs for the business overheads is determined.

A business overhead mark-up of 1.97%³⁶ is used for the capital expenditure costs.

9.9 Output results

The final LRIC costs are generated for each service using a multiplication factor to convert the LRIC cost in minutes/lines per annum into an appropriate cost – either:

- AUD/line/month (access line services)
- AUD cents/minute (voice traffic services)
- AUD cents/call (LCS service)
- AUD/Mbit/s (other data transmission services).

The Core network results are presented as a marked-up LRIC cost for each of the modelled services.

The Access network results are presented as a Band 1, Band 2, Band 3/4 (clustered) and Band 3/4 (spread) access cost for each of the appropriate services.

The costs that are attributed to the other core platform costs (other platforms, other fibre services and other duct services) are also output from the Cost module.

³⁶ Uplift % = business overheads attributed to network / network costs.

Annex A: Deriving the Location and Demand database

The Location and Demand database is a set of nationwide locations that Analysys has derived. They are the proxy of the locations of demand that are used for the sampled ESAs. The database contains a set of 9.76 million co-ordinates with an associated level of demand for each set of co-ordinates.

This section describes, in more detail, how this set of locations was derived, using the G-NAF and several additional data sets and is set out as follows:

- Section A.1 describes the G-NAF database and its shortcomings as a location database, in terms of its location data and its ability to identify business addresses.
- Section A.2 describes how we have attempted to clean up the data from the G-NAF database.
- Section A.3 describes how demand was mapped onto this subset of entries and was then converted into our Location and Demand database.

A.1 Background to G-NAF

Analysys has been provided with the May 2007 release of the G-NAF for Australia. This data set is a street address index for the entire country, with separate files for the nine states and territories of Australia:

- Australian Capital Territory (ACT)
- New South Wales (NSW)
- Queensland (QLD)
- Victoria (VIC)
- Northern Territories (NT)
- Western Australia (WA)
- Southern Australia (SA)
- Tasmania (TAS)
- Other Territories (OT).

The data has been compiled from three main sources: the mapping agencies and land registries of each of the Commonwealth, state and territory governments; the Australia Post; and the Australian Electoral Commission.

Each address is geocoded to a particular point on the Earth's surface (using a latitude and longitude). The database also includes other location property information in its tables, from which certain properties of its location can be discerned. The fields included within the G-NAF are shown below in Figure A.1.

<i>Field</i>	<i>Description</i>
G-NAF_PID	Primary identifier
SEARCH_ADDRESS	Full street address (multi-dwelling buildings have a single value here)
ADDRESS_TYPE	Type of address (urban, rural, unknown, house, etc.)
FLAT_NUMBER_PREFIX	Characters prefixing the flat number (e.g. GF)
FLAT_NUMBER	Flat number in a multi-dwelling building
FLAT_NUMBER_SUFFIX	Characters following the flat number (e.g. U, A, etc.)
FLAT_TYPE	Specification of the type of separately identifiable portion within a building or complex
LEVEL_TYPE	Level type
LEVEL_NUMBER_PREFIX	Rarely used
LEVEL_NUMBER	Indicates the floor of the address (0 implies the ground floor)
LEVEL_NUMBER_SUFFIX	Rarely used
BUILDING_NAME	Combines both building / property name fields
NUMBER_FIRST_PREFIX	Rarely used: entries corresponding to a lot of land have the value "LOT"
NUMBER_FIRST	First street number in the range (single dwelling buildings use that in SEARCH_ADDRESS)
NUMBER_FIRST_SUFFIX	Rarely used
NUMBER_LAST_PREFIX	Rarely used
NUMBER_LAST	Last street number in the range
NUMBER_LAST_SUFFIX	Rarely used
STREET_NAME	Street name
STREET_TYPE	Street type
STREET_SUFFIX	Rarely used
LOCALITY_NAME	Locality name (not provided in all cases)
STATE_ABBREVIATION	State abbreviation
CONFIDENCE	Enumerates how many of the three main data sets contain this address
RELIABILITY	Spatial precision of the geocode expressed as a number in the range 1–6
ALIAS_PRINCIPAL	A = Alias record; P = Principal record
LEVEL_GEOCODED_CODE	Indicator of the degree of geocoding this address, from 0–7
LONGITUDE	Longitude coordinate in decimal degrees using the GDA94 system
LATITUDE	Latitude coordinate in decimal degrees using the GDA94 system

Figure A.1: Main fields in the G-NAF [Source: Analysys]

The field G-NAF_PID should be used as the unique identifier of each G-NAF entry: the coordinates and SEARCH_ADDRESS cannot be used since different entries can be geocoded to the same coordinates. In particular, this has meant that this *address* database needs to be processed before it can be used as *location* database.

For the purposes of our database, we would like to discern locations of demand and, in addition, identify where inhomogeneous demand would occur from businesses and multi-dwellings. G-NAF is unable to do either of these without additional processing.

A.1.1 Location data within G–NAF

The G–NAF is constantly being improved in its coverage of addresses and accuracy of location data. The accuracy of entries is captured by several fields in the G–NAF, but the two most important are *confidence* and *reliability*:

- Confidence is determined by how many of the three primary sources can validate the address and takes values between 0 and 2. A value of 2 means that it can be found in all three primary sources, whereas 0 means that the entry can only be identified in one of them.
- Reliability refers to how accurately the entry has been geocoded. In the current issue of G–NAF, this can take values of 2–6,³⁷ with:
 - 2 or 3 indicating that an address has been geocoded to within the cadastre for that address
 - 4 indicating geocoding to the proximity of a unique road feature, whereby the position is given as the centre of that road
 - 5 indicating geocoding to within a locality or neighbourhood, whereby the position is given as the centroid of that region
 - 6 indicating geocoding is only to within a region, whereby the position is given as the centre of that region.

As a result, entries with a reliability of 4 or more may be co-located with several entries despite being at different locations in reality. The number of entries in G–NAF by territory and reliability is given below in Figure A.2.

<i>State</i>	<i>Reliability 2 or 3</i>	<i>Reliability 4</i>	<i>Reliability 5</i>	<i>Reliability 6</i>	<i>Total</i>
ACT	174 499	6661	121	28	181 309
NSW	3 586 106	503 009	52 877	8	4 142 000
NT	77 346	9556	339	744	87 985
OT	578	640	0	474	1692
QLD	2 356 195	305 247	23 918	86	2 685 446
SA	838 176	91 328	25 859	84	955 447
TAS	272 297	27 629	2116	5	302 047
VIC	2 779 095	175 276	20 939	0	2 975 310
WA	1 124 362	128 136	5288	0	1 257 786
Total	11 208 654	1 247 482	131 457	1429	12 589 022

Figure A.2: Number of entries in G–NAF by state and reliability [Source: Analysis]

³⁷

A value of 1 means that the location has been located to a surveying standard, but inspection of this issue of the G–NAF indicates that no entries take this value.

Cumulatively, 89.04% of addresses have a reliability of 1–3, 98.94% have a reliability of 1–4 and 99.99% have a reliability of 1–5. Conversely, 10% of addresses cannot be geocoded to their own land parcel, which means that the coordinates provided by G–NAF in these cases will not accurately reflect the locations demand in the ESA.

A.1.2 Identification of businesses

Identifying business addresses does not appear to be possible using the G–NAF on its own. Candidates for such identifiers all appear to have severe limitations:

- ADDRESS_TYPE does not have an explicit identifier for businesses and the field appears to be sparsely populated except for the value such as ‘Unknown.’
- FLAT_TYPE has categories that can be directly associated with business types (e.g. office, shop, factory, etc.). However, there are too few of these addresses to cover all business sites in Australia.
- ADDRESS_TYPE sometimes uses the value “Unit,” which could be seen as an indicator of a business site. However, an Internet search of a selection of these addresses indicates that they can be either residential or business addresses.

A.2 Pre-processing the G–NAF

The G–NAF database was designed to be a database of addresses rather than locations in Australia. As a result, a series of pre-processing steps have been required to clean up the database for the purposes of the geoanalysis and access network module. Various sources are used in this process, as summarised below in Figure A.3.

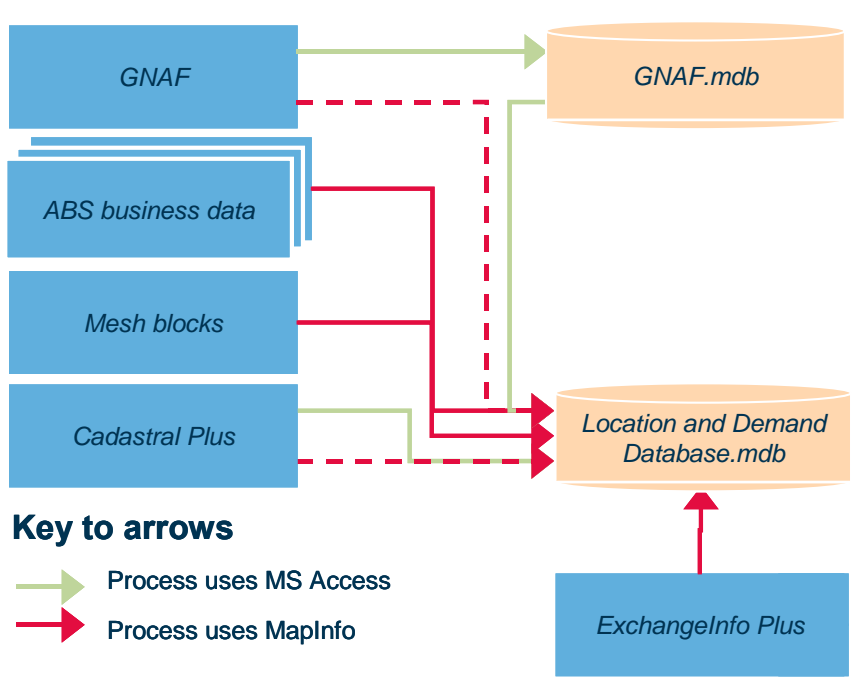


Figure A.3: Processing required to filter the G-NAF [Source: Analysys]

Every G-NAF entry was first allocated to an ESA using ExchangeInfo and a land use using the mesh blocks. The mesh blocks split Australia into approximately 300 000 regions, with particular granularity in urban areas. Each is designated as one of nine different land uses, which we have assumed to either be residential or non-residential, as shown below in Figure A.4.

Land use	Residential or non-residential
Residential	Residential
Agricultural	Residential
Parkland	Residential
Other	Residential
Commercial	Non-residential
Education	Non-residential
Industrial	Non-residential
Hospital/medical	Non-residential
Transport	Non-residential

Figure A.4: Land uses given by mesh blocks [Source: Analysys]

Whether a G-NAF entry is residential or non-residential in nature is inferred by this mapping.

As explained previously, the G-NAF needs to be cleaned due to the limitations of parts of the location data. It is currently updated and maintained quarterly and this process leads to new records being inserted and existing records being updated or even retired. It contains approximately 12.6 million entries in all, but the Australian Bureau of Statistics (ABS) indicates far fewer residences and registered businesses (as of mid-2006) than this, as shown in Figure A.5.

State	Data from the ABS			G-NAF
	Households	Registered businesses	Addresses ³⁹	Total entries
ACT	128 500	10 184	138 684	181 309
NSW	2 569 700	285 011	2 854 711	4 142 000
NT	53 200	52 88	58 488	87 985
OT				1692
QLD	1 545 000	157 134	1 702 134	2 685 446
SA	637 900	53 034	690 934	955 447
TAS	200 200	15 402	215 602	302 047
VIC	1 987 900	204 295	2 192 195	2 975 310
WA	803 700	77 505	881 205	1 257 786
Total	7 926 100	807 853	8 733 953	12 589 022

Figure A.5:
Households by state,
compared with total
G-NAFs [Source:
ABS,³⁸ G-NAF]

Clearly, G-NAF implies far more addresses than the ABS data. The reasons for this include that:

- G-NAF is built from three sources, so addresses often can come under two or more different labels (for example, if the building lies on the corner of two streets).
- Inaccurate or outdated data from the primary sources will follow through into the G-NAF.

Therefore, reasonable ways have been identified in which to remove entries from the G-NAF to try and reduce this excess, without destroying the good location data that it does possess. As a starting point, all of the addresses with a reliability of 4 or more have been deleted. These points are not properly geocoded to a land parcel, so the locations are the most inaccurate of all. In addition, entries have been deleted where they either:

- are labelled as alias addresses
- are not lying within a land parcel in the Australian cadastre
- have a confidence of 0 and a reliability of 3 (these are addresses that only appear in one primary source and can only be placed near the land parcel, or *possibly* within it)
- are labelled as lots (see Figure A.6), but share a land parcel with at least one other entry that was not labelled as a lot.

9.8 million addresses survived these tests, which are explained in more detail below.

³⁸ The data set used for businesses is from the ABS, 8165.0 - Counts of Australian Businesses, including Entries and Exits, Jun 2003 to Jun 2006. This data set provides the number of registered businesses for each postal area in Australia.

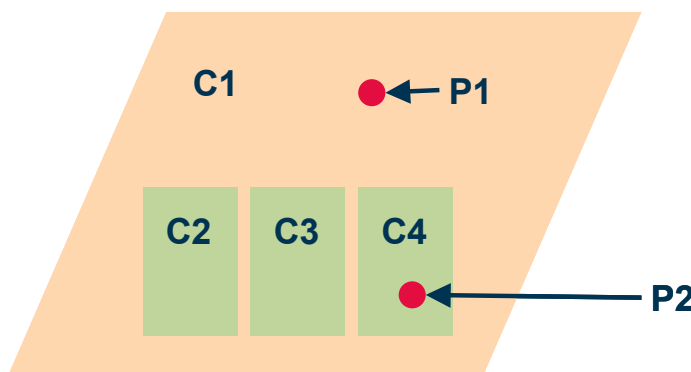
³⁹ Does not include businesses with no employees.

Alias addresses

Addresses are normally identified by a single official, or *principal* address. However, addresses can also have unofficial labels that may also be used to identify the address and are included for the purposes of address searching. An example of this is for a household on the corner of a street, which can be referred to by other street name. These unofficial addresses are also included in the G-NAF as *alias* addresses. Our emphasis is on the demand at individual locations in Australia rather than address searching. Therefore, Analysys has removed all of these alias addresses.

G-NAF entries lying outside land parcels

In order to complete this sifting process, each G-NAF entry was mapped to a unique land parcel using CadastralPlus. In some cases, land parcels overlapped with each other. In these cases, the G-NAF entry was mapped to the smaller land parcel. This is because it was often the case that an overlap would occur with a more granular parcel overlaid on top of a larger one, as demonstrated below in Figure A.6.



*Figure A.6:
Demonstration of
overlapping land parcels.
In this case, P1 is
allocated to C1 and P2 to
C4 [Source: Analysys]*

Confidence of 0 and reliability of 3

Using the definitions of these two measures, these are interpreted to be addresses that only appear in one primary source and can only be placed near the land parcel, or *possibly* within it. The combination of these values for these two measures means that the geocoding of such entries is likely to be of lower quality.

Lots

A lot is a land parcel on which several plots (for individual dwellings) can be built. Therefore, this entry is usually in addition to an entry for each address, meaning that there will be an excess of one G-NAF entry in any land parcel in these cases. However, in some cases, a land parcel would contain the entry labelled as a lot as its only entry, in which case it would want to be retained, so

that demand to that land parcel would be recognised in the geanalysis and access network module.

The sifting does lead to excessive removal of entries from rural ESAs due to their poorer level of geocoding. To try and redress this, sifting is not permitted to occur wherever an ESA would lose more than a certain proportion of its G-NAF entries. This additional step leads us to a final *adjusted* set of addresses.

A.3 Demand mapping

The Location and Demand database requires a demand associated to each location. We have assigned demand to G-NAF entries first and then grouped entries by location, aggregating the demand. A plausible distribution of demand at the address level has that low-demand addresses are predominantly in residential areas and high-demand addresses are in non-residential areas. This is separate to the high demand of multi-dwellings, which are captured through the aggregation stage. The process implemented to assign demand is summarised below in Figure A.7.

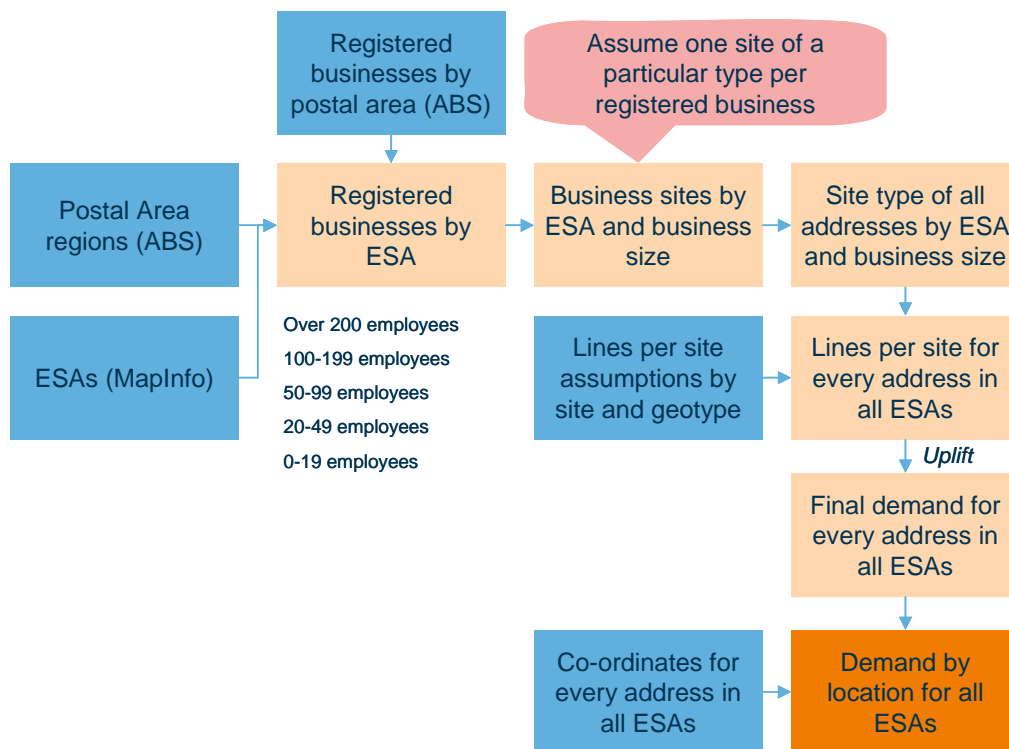


Figure A.7: Process for deriving inhomogeneous demand at locations [Source: Analysys]

The data sets used in the modelling are as follows:

- surviving G-NAF entries, classified as either residential or non-residential using mesh blocks
- registered businesses in Australia, by both postal area and size from the ABS
- postal area boundaries, of which there approximately 2500, again from the ABS Web site

- ESA boundaries from ExchangeInfo.

For our purposes, we have ignored data of businesses which have no employees (i.e. self-employed), since these are likely to be home-operated businesses, whose demand in terms of lines required is unlikely to be distinguishable from residential services. Moreover, according to ABS data on businesses,⁴⁰ there are very few (less than 50 000) such businesses nationwide that have their own separate premises.

Identifying registered businesses by ESA and defining site types

We have used MapInfo and ExchangeInfo to map these registered businesses by postal area onto ESAs. We have then defined each business size to have its own site type, with associated demand. Hence, a registered business in an ESA with over 100–199 employees will have a “100–199” site, which requires a different demand to a “20–49” site. The site types used are 1–19, 20–49, 50–99, 100–199 and 200+. We have merged the 1–4 and 5–19 business sizes since these are unlikely to have different site types, both being effectively SMEs.

Our assumption is then that each registered business of a particular type in an ESA will have a site of that type located within the ESA. We acknowledge that using registered businesses does not equal business locations where a business has multiple sites. This is a limitation but a reasonable attempt to capture the size of businesses given the available data.

Mapping sites to G–NAF entries

For each ESA, we have then used Visual Basic, to identify, at random, each of these sites with a G–NAF entry, using the following (replicable) algorithm:

- For each ESA, we record how many of each site type are estimated in the previous stage.
- Starting with the non-residential G–NAF entries and the largest site type which is calculated to be needed for the ESA, we randomly pick an entry and state it to be that site type.
- Continue randomly picking entries and allocating them to this site type until we have enough of this site type.
- Repeat for the other site types calculated to exist in this ESA, in descending order of size (200+, 100–199, 50–99, 20–49 and 1–19).
- Any remaining non-residential addresses are labelled as non-residential.
- Once an entry has been allocated a site type, it cannot be allocated a different site type.

⁴⁰

8127.0 - *Characteristics of Small Businesses*, Australia (Reissue), 2004, Australian Bureau of Statistics.

- If there are insufficient non-residential entries to allocate to sites, then residential entries are allocated at random once the non-residential entries are exhausted.
- If there are insufficient entries in total, then non-residential entries are allocated sites first, followed by residential entries. The remaining sites are left unallocated, but should be smaller sites. During this process, this was found to happen on ESAs with a very small number of G-NAF entries.

The benefit of allocating sites in descending order is that, should there not be enough non-residential sites, only the smaller sites, which presumably have smaller line requirements, appear in residential areas. This is less of an issue in terms of unrealistic demand.

The type of output that this algorithm can generate is demonstrated below in Figure A.8.

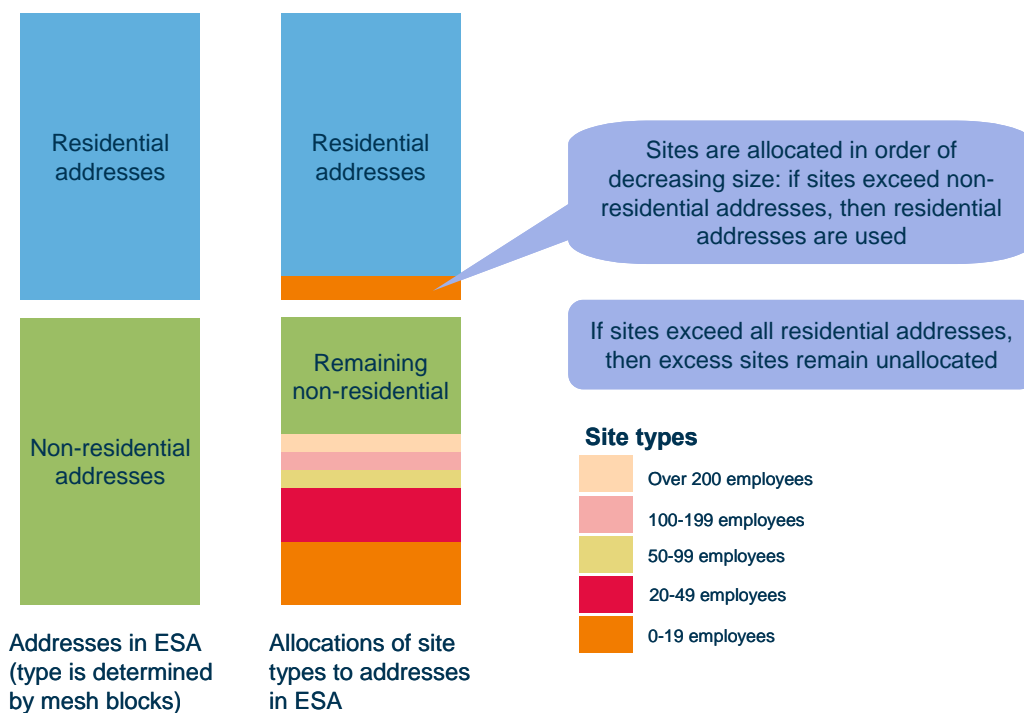


Figure A.8: Diagram of output from allocating business site types to G-NAF entries [Source: Analysys]

We note that the larger registered businesses can have several sites, but these are usually smaller sites and we assume that these lie across Australia at the non-residential G-NAF entries that do not have a particular site allocated to them, remaining classified as “non-residential.”

Mapping demand to sites

For each G-NAF entry, we have the site type for each location. Two further calculations are then completed on the database: a basic demand mapping followed by an uplift.

► *Basic demand mapping*

We have assumed that each site type has its own line requirement. This is a set of parameters that we currently vary by ULLS band. These parameters can be varied in the Location and Demand database using a sequence of queries. These parameters could be varied by geotype if so desired. The values currently assumed are shown below in Figure A.9.

<i>Site type</i>	<i>Band 1</i>	<i>Band 2</i>	<i>Bands 3/4</i>
Residential	1	1	1
Non-residential	2	2	1
0–19 employee site	3	2	1
20–49 employee site	6	2	2
50–99 employee site	12	4	4
100–199 employee site	20	6	6
Over 200 employee site	75	22	22

Figure A.9: Line requirement assumptions by ULLS band [Source: Analysys]

These line requirements are based on estimates of the ratio of lines per site for businesses that Analysys has sourced from previous projects in other jurisdictions. In Band 1, we assume that residential sites require 1 line and non-residential sites require 2 lines. The number of lines of the other site types are scaled up according to these estimated ratios. For example, our estimates indicate that the largest business size has 3.75 times the lines per site of the “100–199” businesses.

We assume lower lines per site in the other bands:

- In Band 2, we also assume that the 0–19 and 20–49 businesses also have 2 lines and only scale the remaining three site types using the appropriate ratios.
- In Bands 3 and 4, we assume that, whilst the 20–49 businesses have 2 lines per site as in Band 2, we assume that the 0–19 businesses only have 1 line per site.

Each G–NAF entry is then assigned the appropriate demand given its ULLS band and site type.

► *Uplift*

Analysys has been provided with the number of services in operation (SIOs) on an ESA basis.⁴¹ This allows us to calibrate our modelled levels of demand on a fairly granular level. Having completed the basic demand mapping, variations across a geotype mean that some ESAs have insufficient demand compared to actual levels, whilst others have an excess demand. The latter is entirely possible given the other fixed networks in Australia, which have taken customers off of the incumbent. However, the former is unrealistic. To try and reduce any such deficits, in particular ESAs, incremental units of demand have been added to G–NAF entries according to the following algorithm:

⁴¹

Data is part of RKR requirements. Report is dated September 2007. SIOs include PSTN, ULLS and xDSL only.

- Any ESAs with fewer than a certain number of SIOs are not adjusted further, since the differences are unlikely to be consequential.
- Any ESAs with more units of demand than actual SIOs are not altered either.
- Remaining ESAs, which have fewer units of demand than SIOs, are considered in turn.
- The deficit between the current levels of demand and actual incumbent SIOs for the ESA is calculated.
- Single units of demand are added at random to the residential G–NAF entries.
- Once a residential entry has received an extra unit of demand: it may not receive another until all other residential entries in that ESA have received an extra unit of demand.
- This continues until the number of units of demand is equal to the number of incumbent SIOs in that ESA, or a maximum limit of additional SIOs has been added to all residential addresses.
- Should the residential addresses be filled and there are still fewer SIOs than units of demand, then the same process is repeated on all non-residential entries in the ESA, up to a maximum limit of additional SIOs on each non-residential address.
- If there is still a deficit after both residential and non-residential G–NAF entries have all been fully uplifted, then no further action is taken.

The current assumptions used in this algorithm are:

- An ESA must have at least 100 Telstra SIOs, according to the RKR data, for it to be considered by the uplift algorithm. This assumption is used to reduce calculation times: it is estimated that the discrepancy is approximately 0.3% of the total target demand base.
- A maximum of 2 additional units of demand to be added to a residential address.
- A maximum of 5 additional units of demand to be added to a non-residential address.

These outputs give a total of 11.3 million units of demand across Australia, compared with 10.3 million actual SIOs utilised by the incumbent.⁴² An excess is expected given the competition of other networks in Australia. However, this number can be calibrated further if required, by changing one of the following elements:

- number of G–NAF entries
- assumptions for lines per site type
- assumptions contained in the uplift algorithm.

⁴² Data is part of RKR requirements. Report is dated September 2007. SIOs includes PSTN, ULLS and xDSL only.

Deriving locations

In order to group G-NAF entries to locations, we do not group by land parcel. This is because:

- The identifier used for land parcels in CadastralPlus is not always unique, so grouping by it could lead to aggregation across several parcels.
- Certain land parcels, especially in rural areas, contain several (or even dozens of) distinct households. Such cases would lead to erroneous aggregation.

G-NAF entries have therefore been grouped by the latitude and longitude of their G-NAF. Since these are decimals, Microsoft Access cannot group these accurately. Therefore, the entries have been grouped using the nearest whole number of the original longitudes/latitudes multiplied by 50 000. This is equivalent to grouping any coordinates that lie within approximately 3m of each other. Demand is aggregated to locations. Given that the grouped co-ordinates may differ very slightly, the average values at a locations are calculated.

Annex B: Designing the sample of ESAs

B.1 Sample design

There are numerous ways to derive sample sizes for each stratum (i.e. geotype) in a population given a total sample size. Analysys has based its derivation on a methodology called optimum allocation.⁴³ This states that, for a sample of size N across M strata (each of size n_k), the sample size h_k for stratum k should be:

$$h_k = \frac{Nn_k}{\sum_{i=1}^M n_i \sigma_i}$$

where σ_i is the standard deviation of the variable being investigated by the sample, for the k th geotype. In our case, the variable we are effectively investigating with our sample is the cost per line in Australia.

Although we do not know the standard deviation of the cost per line for each geotype, we have identified a suitable variable strongly related to the cost per line which can be calculated for all ESAs, so that its standard deviation across each geotype can be used as a reasonable proxy.

In order to use this formula, Analysys has made two important assumptions. When calculating the sample size for the k th geotype:

- n_k is interpreted to be the total SIOs within a geotype, rather than the number of ESAs.
- The value of σ_k is estimated using the standard deviation of road density per SIO, calculated across all ESAs within the geotype.

Using the number of ESAs in a geotype for n_k is the obvious choice but leads to almost the entire sample being dedicated to geotypes 12–14, with these three geotypes having many ESAs with little demand. It is more appropriate for the sample to adequately focus on the costs associated with the most SIOs, which lie in Bands 1 and 2.

There are no previous samples of our geotypes on which to estimate the standard deviation of the associated cost of service per line. We have used the standard deviation of the road density per SIO as a proxy, given that the cost of service for a SIO is driven by the trench needed to lay cable to the location. Using road *density* per SIO normalises for ESAs with large areas, sparse population and large amounts of trans-national road.

⁴³ Thompson, S., *Sampling*, Wiley–Interscience Publication, (US, 1992).

Applying the formula in this manner leads to the sample summarised below in Figure B.1.

<i>Geotype</i>	<i>Band(s)</i>	<i>ESAs</i>	<i>SIOs</i>	<i>S.d. of road density per SIO</i>	<i>Sample size</i>
1	Band 1	3	44 804	0.00020051	3
2	Band 1	14	237 009	0.00116018	4
3	Band 2	139	1 909 188	0.00063425	16
4	Band 2	219	2 825 812	0.00061397	24
5	Band 2	160	1 747 611	0.00126436	30
6	Band 2	66	494 792	0.00102969	7
7	Bands 3/4	91	224 760	0.00247458	8
8	Bands 3/4	107	153 554	0.00263271	6
9	Bands 3/4	86	41 597	0.00556077	3
10	Bands 3/4	181	832 708	0.00183476	21
11	Bands 3/4	236	608 141	0.00088575	7
12	Bands 3/4	844	728 543	0.00222076	22
13	Bands 3/4	1936	395 470	0.00685165	37
14	Bands 3/4	912	67 862	0.01305765	12
Total		4994	10 311 851		200

Figure B.1: Summary of sample [Source: Analysys]

The formula determines that the first geotype should not be represented by the sample, since it only has three ESAs. Given its small size, we have elected to simply do a census.

B.2 Identifying the sampled ESAs

Although, the sampling process should create a statistically random sample, it should nonetheless minimise the number of ESAs that require processing for the geoanalysis and access network module in the event of a revised sample: for example, if the sample size or geotypes were revised. This is because converting the location data is time-consuming, so being able to retain processed ESAs in a revised sample saves processing time.

The ESAs were ordered and given a permanent index of 1 to 5070. A simple piece of Visual Basic was used to generate, randomly ordered lists of the first 10 000 integers, with at least one for each geotype. 25 lists were generated should the number of geotypes ever be increased or ESAs were segregated into sub-areas. The ESAs were then arranged in ascending order, first by geotype and then by index. For the k th geotype requiring n ESAs in the sample, the first n ESAs of that geotype in this ordered list corresponding to the first n indices in the k th list of integers were selected.

Suppose the sample for the second geotype was reduced from seven to five ESAs. In this case, we would simply remove the sixth and seventh ESAs of this geotype. Conversely, if the sample increases to nine, with no changes in the geotypes, then we should be able to use the first seven ESAs already processed, meaning that only the eighth and ninth ESAs need to be prepared.

The ESAs sampled in Bands 1 and 2 as a result of this selection process are shown below.

Geotype	ESA	Geotype	ESA	Geotype	ESA	Geotype	ESA
1	WAYM	3	SMRN	4	SALA	5	GOSN
1	BWER	3	COOG	4	EMPS	5	CRCF
1	FLNF	4	AARE	4	LSDE	5	CTON
2	EXHN	4	TUTT	4	WOBB	5	NBUR
2	SGHL	4	NGLG	4	CURN	5	SEFD
2	CHLT	4	FKNR	5	SPHS	5	GRMT
2	RASH	4	KELL	5	TQAY	5	NDLN
3	REDF	4	SHAL	5	IPSW	5	CPBA
3	EDWN	4	MYDE	5	MDLE	5	DKIN
3	FTON	4	DRRA	5	DDLY	5	SUTH
3	CMLL	4	SOTH	5	CURW	5	LAVN
3	CTAM	4	CLVY	5	NLBY	5	HOME
3	MIRA	4	MTHN	5	CTOE	5	BSDN
3	HTLL	4	WRWD	5	MINT	5	BEEL
3	GLIS	4	CSEA	5	COWS	6	LINE
3	REDC	4	TAOA	5	FMTL	6	BKPK
3	NRWD	4	STLE	5	BBRY	6	ORHS
3	MCHN	4	BLHS	5	SYRA	6	MAIT
3	CPRO	4	WLSD	5	SATL	6	BLNA
3	JVLE	4	FREN	5	BEND	6	BOWL
3	AVAL	4	GUGA	5	BAXR	6	BLCT
3	AVAL	4	GUGA	5	BAXR	6	BLCT

Figure B.2: List of sampled ESAs for Bands 1 and 2 (geotypes 1–6) [Source: Analysys]

Similarly, the sampled ESAs in Bands 3 and 4 are shown below in Figure B.3.

Geotype	ESA	Geotype	ESA	Geotype	ESA	Geotype	ESA
7	SHHD	10	DROU	12	MTTA	13	YRCK
7	WONE	10	PMWS	12	FELU	13	CRRL
7	LUCI	10	GLFD	12	UPWI	13	ELBH
7	BNHD	10	THIR	12	MLNG	13	MTLE
7	SPWD	10	YMBA	12	MRLD	13	DTWY
7	RONE	10	WOOCR	12	DGPK	13	LURA
7	RABH	10	BATE	12	BULS	13	KBLA
7	MONV	10	BRIL	12	BGHT	13	FLFT
8	ROUS	10	BLIB	12	WUND	13	HUMU
8	MNBH	11	ADLE	13	PENR	13	HKTN
8	MRNB	11	MOSS	13	UORA	13	GORK
8	BUND	11	OAKS	13	DIRR	13	TGHA
8	WELL	11	WYRE	13	GRBT	13	BODA
8	MVIC	11	LLDO	13	PPNG	13	GNLN
9	MTLD	11	PITN	13	MYRE	13	BAEN
9	MMHK	11	DUNX	13	TBRI	13	CRAC
9	TALB	12	MTMO	13	REIR	13	BWLK
10	GOBH	12	MARY	13	YAMS	14	YEUN
10	LOWO	12	TARV	13	LLAN	14	TMOR
10	SANR	12	WAMP	13	BINO	14	LORQ
10	DBGH	12	CRBG	13	BEEH	14	QDAN
10	LITH	12	BLNG	13	HECK	14	TANE
10	OCNS	12	MBGH	13	CNHM	14	MATG
10	NING	12	KEWA	13	MUNR	14	HOLT
10	MONB	12	KTHN	13	WOMT	14	WEDG
10	CASM	12	HODD	13	MUTA	14	LARM
10	YEPP	12	KATG	13	LBID	14	MNOA
10	NRDS	12	JAPO	13	SWAM	14	FLNS
10	COLA	12	NATU	13	WDYR	14	CRCE

Figure B.3: List of sampled ESAs for Bands 3 and 4 (geotypes 7–14) [Source: Analysys]

Annex C: Clustering and spanning tree algorithms

This annex provides further details on the clustering and spanning tree algorithms outlined in section 5.2.

C.1 Clustering algorithms

The clustering algorithms take a set of locations and groups them subject to a capacity and a distance criterion. The algorithm has two phases: a creation phase and a refinement phase.

C.1.1 Creation phase

This groups the points into a set of clusters, beginning with the full set of points. The demand-weighted centre is calculated for this “parent cluster.” “Child” clusters are grown by adding points from the parent cluster incrementally. The point that is furthest from the parent cluster’s centre is used as the seed for a new child cluster. Points that would not exceed the child’s cluster capacity and would not break the distance criterion are then added incrementally until no more addition is possible. Each time a point is added to the child cluster, the demand-weighted centres of both the current child and parent clusters are re-calculated. This continues until the parent cluster itself satisfies both criteria. This is known as top-down clustering and fixes the number of clusters.

C.1.2 Refinement phase

The creation phase does not necessarily create an optimal set of clusters. A collection of refinement algorithms have been designed to improve the clusters, as described below:

Simple re-assignment

For each point in turn, the cluster identifies the cluster (after itself) whose demand-weighted centre is closest to its own. The point is then moved to this cluster if the following conditions are satisfied:

- the point is closer to this demand-weighted centre than its own
- the new cluster has sufficient spare capacity
- all points in the new cluster obey the distance constraint with respect to the re-calculated DP cluster centre.

This process continues cycling through all points, multiple times if necessary, until no more re-assignment is possible.

Full optimisation

For each cluster, the total distance between all the points and the cluster centre is derived. Then, it cycles through all points in turn, multiple times if

necessary, until no points have been moved in a whole cycle. The loop:

- Identifies the cluster of the point and the total distance (d_1) between all points in this cluster and its cluster centre.
- Temporarily removes the point from its cluster and re-calculates the demand-weighted cluster centre and the total distance (d_2) between the cluster points and the new cluster centre.
- Restores the point to its cluster and the demand-weighted centre.
- For each cluster with spare capacity, stores the total distance (d_3) between all points in this cluster and its current cluster centre.
- Adds the current point into the cluster and re-calculates the demand-weighted cluster centre and the total distance (d_4) between the cluster points and the new cluster centre.
- Finds the cluster which gives the largest reduction in total distance (i.e. which maximises $[(d_1-d_2)-(d_4-d_3)]$).
- The point is moved to this cluster if it maximises $[(d_1-d_2)-(d_4-d_3)]$ and would also satisfy the normal distance constraint using its new demand-weighted centre.

Swap

For each point in turn, this algorithm:

- Identifies the cluster whose demand-weighted centre is closest.
- If the cluster is not its current cluster and, if moving the point to the new cluster violates the capacity criterion, then tries to find a point in the new cluster which can be swapped with the original point so that:
 - both new clusters satisfy the cluster capacity constraint
 - the sum of the two distances between the points and the old cluster centres is improved compared with before
 - if the first two are true, then it is also checked if the sum of the two distances between the points in their original clusters and their original cluster centres is less than the sum of the two distances between the points in their new clusters and their new cluster centres
 - both obey the distance constraint with respect to the new

demand-weighted centre of their new clusters.

- If such a point is found, then it revises the two clusters and re-calculates their demand-weighted centres.
- If several such points are found in the new cluster, then it uses the point which reduces the sum of the distances by the most.

High-demand swap Same as *swap*, except that it only considers points where demand is greater than 1 and will consider swapping them provided the two new cluster capacities are both less than a higher capacity constraint (see below).

High-demand simple re-assignment Same as *simple reassignment*, except that it only considers points where demand is greater than 1 and will consider re-assigning them provided the cluster which it is moved to does not exceed a higher capacity constraint (see below) and the distance criterion.

Single-demand swap Same as *swap*, except that it only considers points with a demand of 1: the two new clusters must automatically satisfy the capacity constraint, but the distance criterion must still be satisfied.

Consolidation For pillar clusters, an additional algorithm has been designed to try and merge small clusters into other clusters. The pillars are ordered in ascending capacity and all possible pairs are cycled through. If two pillars can be merged, with the merged cluster satisfying the higher capacity constraint and distance criteria, then they are merged. This continues until no further consolidation is possible.

In the above, we refer to a “higher” capacity constraint. This is because an access network is not designed assuming that its full capacity will be used immediately upon deployment. This is to account for spares and for future additional demand. For example, if a pillar has a “designed” capacity of 400 lines, then it may be deployed to serve a proportion of this (for example, a 10% spare capacity means that 360 lines are used upon deployment).

The consolidation process never exceeds the absolute pillar capacity (400 pairs), but it may lead to the practical pillar capacity being exceeded (i.e. some of the spare capacity is actually used).

C.1.3 Additional assumptions

There are several other assumptions used in these processes that are worth mentioning, namely:

- location and treatment of the RAU
- prohibition of moving the RAU during cluster refinement

- using different sequences of clustering refinement algorithms for clustering of different network elements.
- implementation of a simplified distance metric
- scaling of demand for the purposes of wireless clustering.

Location of RAU The RAU in each sampled ESA is currently assumed to be the first point in the list of locations (which is the location of the actual copper centre in that ESA, according to ExchangeInfo), using a flag on the worksheet for the sampled ESA. This point is given a nominal demand of one unit so that it can be included in the spanning tree algorithm.

However, if the flag is left blank, then the algorithm will calculate its own RAU. In such cases, it is calculated twice:

- prior to pillar clustering, the RAU is assumed to be at the demand-weighted centre of all of the points in the ESA
- following pillar clustering, the RAU is revised to be the demand-weighted centre of all the pillars in the ESA.

Prohibited removal of the RAU between clusters in ESAs using fibre rings The clustering algorithms are employed several times during a calculation, using different inputs. In particular, they are used to group pillar nodes into clusters to link with fibre rings. The RAU must lie in each of these clusters of pillar nodes and this is set up at the start. Therefore, during clustering, the RAU cannot be moved between clusters, as the algorithm may end up having the point occurring twice in a cluster. This constraint is employed whenever the algorithms are used for an ESA that is deploying fibre rings.

Variation in refinement phase with use of clustering For each of the occasions where clustering is used in the geoanalysis and access network module, a certain combination of refinement algorithms are used in a particular order to produce the best result. The specific order used for each particular stage is described later in this chapter.

Distance metric used Since clustering does not require the calculation of actual distances, only the comparison of distances (i.e. is the point P1 or the point P2 closer to the point P3), we assume the simpler crow-flies metric for speed of calculation (i.e. $[(x_1-x_2)^2+(y_1-y_2)^2]$).

Scaling of demand for wireless clustering The demand by location stored in the workbooks reflect copper demand in terms of lines. Hence, 1 unit of demand means that 1 copper pair is required. This mapping of demand may not be suitable dimensioning for a wireless solution, as these will be driven more heavily by the Erlangs of traffic passing onto the network.

For example, suppose you have a location that is a small business site. Although the usage is higher than that of a residential user, it can still

conceivably be served by one line in a fixed network. However, if they are served by wireless, then their higher usage will have a greater impact on the dimensioning of this network, which may be appropriate to reflect in the capacity calculations.

This is implemented in the geoanalysis and access network module by applying different scaling factors to locations with one unit of demand compared with those with more than one unit of demand. The module currently has both inputs set to 1, since we assume that this imbalance is not material.

C.2 Spanning tree algorithms

These algorithms take a set of locations and derive a set of edges that form a spanning tree between the locations. The algorithms that we have designed minimise a proxy cost function, which takes the form $k_1*d + k_2*c + k_3*d*c$ for copper or fibre deployment. d is the distance between two points and c is the number of pairs required between the two points (the capacity). The terms of this function estimate the costs of trench, jointing and cable respectively.

We use a modified version of Prim's algorithm to derive the spanning trees within urban DP clusters, rural pillar clusters and to connect wireless nodes back to the copper centre (RAU) using a microwave backhaul network. We use a version of Dijkstra's algorithm to join pillars back to the RAU and into fibre rings, as well as to join urban DPs back to their parent pillars. These are described in more detail below.

C.2.1 Modified Prim

This algorithm begins with the individual points and their respective demands. It starts at the central node. For urban deployments, this corresponds to the DP and for rural deployments this is the pillar. For the purposes of this description, we describe the urban case:

- The algorithm adds vertices to the tree incrementally, by joining an unattached vertex to a vertex in the existing tree using an edge.
- At each stage, the algorithm has stored the identities of all attached points and all unattached points.
- All possible pairs of points containing an attached point and an unattached point are considered.
- For each such pair, the unattached point is temporarily attached to the tree at the attached point and the average cost per unit capacity of the new candidate tree is calculated. The cable size required to serve the unattached point is calculated for the link:

- if the average cost per unit capacity is lower than the previous best value, then the tree is temporarily updated with the required cabling all the way back to the DP.
 - if the average cost per unit capacity is still lower than the previous best value when the cabling requirements have been fully updated all the way back to the DP, then the pair of points is stored as the best pair.
- For the pair of points that generates the lowest average cost per unit capacity overall, the unattached point in the pair is joined to the attached point in the pair.
 - The edge is added to the spanning tree and the lists of (un)attached points are updated.
 - The throughput capacity and cabling at the DP is also calculated.
 - The total network proxy cost of the updated tree is calculated for future comparisons.
 - This process iterates until there are no remaining unattached points.

This algorithm is unable to work effectively with points of zero demand. This has meant that we cannot include street junctions within the set of locations in an attempt to make the algorithm follow the road more closely. Therefore, when viewing the outputs of the modified Prim algorithm in conjunction with a street network, corner-cutting is likely to be observed. However, this is the reason why we have calibrated and subsequently used a p-function to calculate trench/cable lengths accounts, since this uplifts these distances to be longer than crow-flies. This function is described in more detail in section 4.2.5.

C.2.2 Dijkstra

This algorithm begins with a set of cluster centres and a set of points contained within the clusters.

As a starting point, for each pair of clusters, the cheapest pair of points with which to join a pair of clusters through their centres is calculated using the proxy cost function. The points that a cluster centre passes through on the way to linking to the other cluster centres are also stored, so that the total distance traversed along trench between the two points can be calculated.

The subsequent calculation is strictly the Dijkstra algorithm and starts at the central node. For urban deployments, this algorithm is run twice: once to join DPs back to a pillar and once to join the pillars to the RAU. For rural deployments, it is only required for the latter. For the purposes of this description, we describe joining the DPs to their parent pillar. The algorithm then proceeds as follows:

- We assume that initially all DPs are joined directly to the pillar (in a star formation).
- Beginning at the pillar, for every other DP the requirements to link it to the pillar are recalled, in terms of:

- extra trench
- cost of linking the two points in their respective DP clusters (using the proxy cost function)
- cabling cost of linking the two DPs (i.e. excluding trench cost)
- total sheath length between them.

Then, until all DPs in the pillar cluster have been connected to the pillar, we proceed as follows:

- For the given connected DP (the first stage uses the pillar itself), we look through the remaining unconnected DPs and decide which of these can be linked directly to this connected DP with the least proxy cost.
- In order to calculate the proxy cost accurately, trench and cable volumes need to be calculated properly. In the situation shown below in Figure C.4, the trench in the links circled would be double-counted unless the algorithm was refined. We have included a process which only counts this once. It is currently only implemented when connecting pillars to the RAU, where the connections involved would lead to a larger error.




Figure C.4:
Demonstration
situation where a
trench could be
double-counted
[Source: Analysys]

- Having identified the unconnected DP that can be joined to the selected connected DP, we test whether it is a cheaper proxy cost to go from this unconnected DP to the pillar by going through the selected connected DP, or going via the direct connecting path calculated from the unconnected DP to the pillar (i.e. Is the DP connected directly to the pillar, or via existing links?).
- Next, we determine, for each unconnected DP in turn, whether it is more cost effective to get to the pillar via the newly connected DP or remain with its current path. Recall that all unconnected DPs are initially assumed to go directly back to the pillar, but this can evolve with each loop of the algorithm. If it arises that it is more cost-effective, then the path for the disconnected DP should be changed to be via this newly connected DP.

- Finally, we check if it is cheaper for any other DP to be connected to the pillar via this newly connected DP. If it is cheaper, then we change the routing of these DPs.
- We set the selected connected DP to be the newly connected DP and recommence the loop.

This can lead to DPs being daisy-chained back to their respective pillar, or pillars being daisy-chained back to the copper centre.

Annex D: Fibre-ring deployment

Currently, in Band 1, we deploy a fibre ring infrastructure in the access network to connect high-demand locations back to the RAU. This deployment requires two phases. The first phase creates the ring(s) to link the pillars together, whilst the second phase calculates the requirements for the spurs to link the fibre-fed FDPs to the ring.

D.1 Deriving the rings

Deriving the rings uses the following process:

- Identify which pillars are to be included in the fibre ring (the RAU is included by default).
- Calculate how many rings would be needed for each ring to include the RAU and to have less than an upper limit of pillars on the ring.
- Re-calculate the maximum number of nodes that should be on each ring, given the number of pillars and the number of rings. This will allow the rings to be built in more similar sizes. As an example, if at most 20 nodes can lie in a ring, then this assumes that it is preferable for 22 pillars to lie on two rings of 11/12 each, rather than one ring with (for instance) 17 and one with 5.
- Cluster the pillars with this maximum number of nodes as the capacity constraint, and with no distance constraint,⁴⁴ using refinement algorithms in the order set out in the final column in Figure 5.7.
- Use the same algorithm used for joining the pillars to the RAU, but with different proxy cost function coefficients, to identify the best links between any two pillars. This accounts for the fact that some pillar-pillar trenching exists having joined the pillars to the RAU.
- Use the TSP algorithm to identify the best way to join the points in each ring in turn, using the linkages previously costed. Although the TSP algorithm can handle multi-ring problems, we only do one ring at a time here. This is because the algorithm cannot solve a multi-ring problem with more than approximately 25 nodes sufficiently quickly, and this limit is exceeded in some of the largest ESAs, which require many more than 25 pillars.

An example output of for Waymouth is given below in Figure D.1: three rings have been built in this case, using all the deployed pillars. The figure shows the duct deployed for the fibre ring.

⁴⁴ The process has been designed so that the RAU lies in every cluster and cannot be removed from them. This is why the clustering algorithms do not allow the RAU to be moved, so that they can be re-used for the fibre-ring methodology. Not permitting the refinement algorithms to move one point should have negligible effect on the outcome.




Figure D.1:
Fibre ring for
Waymouth, showing
the pillars and the
RAU
[Source: Analysys]

- The algorithm stores the total incremental trench length required to complete the ring structures between the pillars/ LPGS, as well as the total fibre sheath length required.
- The fibre cable length required is calculated assuming that the cable is big enough to carry the entire capacity of the ring demand, using the fibre cable sizes which are presently used. The cable size used on each fibre ring can differ, depending on the total demand of the fibre locations from each pillar on the ring, except for the RAU. The smallest cable with more fibres than this total is then used all around the ring.
- If the total is larger than the largest cable size, then as many multiple sheaths as possible of the maximum cable size are used in order to have enough capacity. However, in order to reduce the excess spare fibre in the ring, if using an extra one of the second largest cables in place of one of the largest size sheaths also gives sufficient capacity, then this is done instead. For example, if the two largest cable sizes are 100 and 300 fibres and the required capacity is 350, then one 100-fibre cable and one 300-fibre cable are used, rather than two 300-fibre cables.

D.2 Linking locations to the ring

When constructing the copper deployment, the algorithms have assumed that a fibre-fed location has a single unit of demand. Therefore, these locations are included in the DP and pillar clusters, and trench is laid to each FDP. A 2-pair cable is also deployed from the DP to the FDP. In order to link the points to the fibre rings, a fibre cable is laid along the same trench that the copper cable traverses in order to get to the pillar. This is assumed to be a discrete cable, which does not merge further up the network until it reaches the pillar. The copper deployments to these locations are assumed to remain in the network as a back-up.

If a fibre ring is not being constructed, then locations are linked to their parent pillar as above. However, the discrete cable is assumed to continue behind the pillar, returning to the RAU via the route taken for the link between that pillar and the RAU.

Annex E: Assumptions in the access network deployment

E.1 Urban deployment

Some of the assumptions that can be used in the urban deployment can vary by geotype. However, for almost all assumptions, the same value is used across all geotypes using the urban deployment. The values used for all such geotypes are shown below in Figure E.1.

Quantity	Value currently assumed
Designed capacity of a distribution point (DP)*	4 units of demand
Capacity of a DP that is reserved as spares	0%
Maximum distance that a point can be from its DP cluster centre*	40-100 metres
Capacity of a pillar	900 pairs
Capacity of a pillar that is reserved as spares	10%
Capacity of cable linking pillars to the RAU	400 pairs
Distribution network cabling utilisation(DP-pillar)	75%
Critical capacity of a pillar	20 units of demand
Main cable size in non-tapered network	100 pairs
Minor cable size in non-tapered network	0 pairs
Maximum distance in distribution network between pits	100 metres
Maximum distance in which cable can be pulled through the distribution network without jointing	500 metres
Maximum distance in which cable between a pillar and the RAU can be pulled through the network without jointing	500 metres
Maximum distance in incremental trench between the pillars allowed between manholes	250 metres
Fibres cable linking a high-demand location to the pillar*	6 fibres
Maximum distance that a DP can be from its pillar cluster centre*	600 metres
Cable sizes used between FDP and DP (pairs)	2, 10, 30, 50, 100
Coefficients for the proxy cost function for the minimum spanning trees within the DP clusters	$K_1=1, k_2=0, k_3=0$
Coefficients for the proxy cost function for joining DPs to their parent pillar using tapering	$K_1=1, k_2=0.1, k_3=0.001$
Coefficients for the proxy cost function for joining DPs to their parent pillar using little/no tapering	$K_1=1, k_2=0.01, k_3=0.00005$
Coefficients for the proxy cost function for joining pillars to the RAU	$K_1=1, k_2=0.1, k_3=0.0001$

Figure E.1: Current assumptions used in the urban deployment (assumptions which vary by geotype are denoted by an asterisk)
[Source: Analysys]

Coefficients for the proxy cost function for joining pillars to each other, for the purposes of building the fibre ring	$K_1=1, k_2=0, k_3=0.001$
Maximum number of nodes in a fibre ring (including the RAU)	20
Minimum units of demand for a location to be served by fibre	40 units of demand
Fibre cable sizes that be employed in the fibre ring (no. of fibres in cable)	18, 24, 36, 60, 120, 312
Distance constraint for installation of an LPGS rather than a pillar (in cable distance)	10 000 metres
Fibre cables linking an LPGS to the RAU	6 fibres
Location of RAU (can be varied by sampled ESA)	Current location for ESA

E.2 Rural deployment

Some of the assumptions that can be used in the rural deployment can vary by geotype. However, all geotypes which use the rural deployment currently use the same assumptions. The values used are shown below in Figure E.2.

<i>Quantity</i>	<i>Value currently assumed</i>
Capacity of a pillar*	900 pairs
Capacity of cable from the pillar to RAU	400 pairs
Critical capacity of a pillar	20 units of demand
Fibre cable linking a high-demand location to the pillar	6 fibres
Maximum distance that a point can be from its pillar cluster centre	4500 metres
Cable sizes used between FDP and the pillar for tapered network*	2, 10, 30, 50, 100, 200, 400, 800, 1200
Main cable size in non-tapered network	100 pairs
Minor cable size in non-tapered network	0 pairs
Maximum distance in which cable can be pulled through the distribution network without jointing	500 metres
Maximum distance in which cable between a pillar and the RAU can be pulled through the network without jointing	500 metres
Maximum distance in incremental trench between the pillars allowed between manholes	250 metres
Coefficients for the proxy cost function for the minimum spanning trees within the pillar clusters with tapering*	$k_1=1, k_2=0.1, k_3=0.001$
Coefficients for the proxy cost function for the minimum spanning trees within the pillar clusters with little/no tapering*	$k_1=1, k_2=0.001, k_3=0.0002$
Coefficients for the proxy cost function for joining pillars to the RAU*	$k_1=1, k_2=0.1, k_3=0.0001$

Figure E.2: Current assumptions used in the rural deployment; (assumptions which do not vary by geotype are denoted by an asterisk) [Source: Analysys]

Coefficients for the proxy cost function for joining BTS back to the RAU*	$k_1=1, k_2=0, k_3=0$
Coefficients for the proxy cost comparison function for fibre backhaul*	$k_1=0.0004, k_2=0.0000062$
Coefficients for the proxy cost comparison function for wireless backhaul*	$k_1=1.6, k_2=1, k_3=5.4$
Maximum distance of a wireless link without using a relay station	30 000 metres
Maximum number of relay stations in a wireless link	5
Capacity of a wireless BTS	250 units of demand
Incremental demand on wireless BTS for single units of demand	1
Incremental demand on wireless BTS for multiple unit of demand, per unit of demand	1
Backhaul capacity per residential subscriber	20kbit/s
Backhaul capacity per business subscriber	20kbit/s
Coverage radius of a wireless BTS	25 000 metres
Coefficients for wireless/copper comparison algorithm	See Figure 5.17
Cost multiplier for a 2Mbit/s wireless link	1
Cost multiplier for an 8Mbit/s wireless link	1.7
Cost multiplier for a 34Mbit/s wireless link	2.3
Minimum units of demand for a location to be served by fibre	40 units of demand
Fibre cable sizes that be employed in the fibre ring (no. of fibres in cable)	18, 24, 36, 60, 120, 312
Fibre cables in link back from LPGS to RAU	6 fibres
Distance constraint for installation of a LPGS rather than a pillar (in cable distance)	10 000 metres
Location of RAU (can be varied by sampled ESA)	Current location for ESA

Annex F: Cable network deployment

Analysys understands that a non-tapered architecture is used in the copper distribution network, i.e. between the pit and the pillar/LPGS. Our access network deployment algorithms, as described in section 5, can deploy cable at this level of the network using either a tapered or a non-tapered architecture: this choice can differ on a geotype basis.

We have processed the sampled ESAs across geotypes 1–14 assuming both a tapered and non-tapered architecture and compared the cost associated with the network volumes required.

Figure F.1 below shows, for each geotype, the total cost of deployment of the trench, cabling, pits and jointing with both architectures.

Band	Geotype	Non-tapered annualised costs (AUDm)		Tapered annualised costs (AUDm)	
		Total ⁴⁵	of which jointing	Total	of which jointing
1	1	14.8	10.1	12.6	8.3
1	2	3.7	1.9	3.5	1.8
2	3	92.0	76.4	99.0	85.1
2	4	139.1	111.2	162.7	138.2
2	5	151.3	106.3	175.8	135.0
2	6	34.4	20.4	39.3	23.8
3 / 4 clustered	7	18.1	10.9	19.0	12.2
3 / 4 clustered	8	7.8	2.6	7.4	2.5
3 / 4 clustered	9	1.0	0.2	0.9	0.2
3 / 4 spread	10	76.6	40.4	79.8	45.5
3 / 4 spread	11	11.7	2.5	12.7	4.2
3 / 4 spread	12	31.3	2.8	30.6	4.0
3 / 4 spread	13	21.4	0.5	18.7	0.4
3 / 4 spread	14	1.8	0.02	1.7	0.01
Total across sample		605.0	386.1	663.5	461.1

Figure F.1: Comparison of tapered and non-tapered cable deployment in our sampled ESAs [Source: Analysys]

As can be seen above, the sampled ESAs in Band 1 can be served more cheaply by a tapered architecture, whilst the reverse is true of the Band 2 ESAs, with each of the four geotypes having a non-tapered architecture as the cheaper alternative. For Bands 3 and 4, the most densely populated geotypes (7 and 10–11) have non-tapered as the lowest cost deployment, whilst the least densely populated geotypes (8–9 and 12–14) have tapered as the lowest cost deployment.

⁴⁵ This cost comprises that of the trench, cabling, pits and jointing within the distribution network.

Across all of the ESAs in our sample, the non-tapered deployment is the lowest cost option.

A tapered architecture is more susceptible to faults, which leads to higher maintenance costs on the jointing. The values above assume that jointing opex is 3% of capex. If we assume a much higher proportion for opex of 15%, then only Band 1 and geotypes 13–14 still have a lower cost tapered deployment.