



Qualities and Constraints in IT Architecture

Performance Engineering

AmGro-From-Home Case Study

Input for Exercise 4 – Estimating System Performance

1. Case Study Introduction – AmGro from Home

1.1 Background to the Case Study

This exercise input material is derived from a real “Home Shopping” business system developed by a major UK food retailer. It is based on the actual business requirements, and closely follows the architecture and design work which was developed as part of the project. However, this is a heavily cut down and short introduction to the project purely for the purposes of this seminar exercise.

As you read this document, remember that your objective is not to produce perfect answers; rather, you will be using the exercise to gain some experience in the performance estimation and modelling techniques presented in the seminar. Therefore the focus should be on *how* you go about the exercises rather than the what particular results you produce in the end.

The questions for this case study are given on slide 47 of the seminar handout. The exercise will either be conducted in-class or given as a small assignment subsequently.

1.2 Project Definition

1.2.1 You ...

You are an Performance Architect in the architecture team recently engaged by Amalgamated Grocers (AmGro) to design and develop their “AmGro-from-Home” Home Shopping system. Your particular responsibilities are in the performance engineering of the solution’s Order Management System. You will be responsible for ensuring the system meets the performance requirements AmGro have set for the new system, and ensuring it will be able to handle the volumes required of it.

You join the project part way through the initial design phase, and the following inputs are already available:

- Business information and goals from interviews with key AmGro personnel, including projected business volumes
- An outline physical architecture for the system, created by the Order Management Lead Architect
- Estimated cost metrics for key transactions

These inputs will be presented later on in this documents.

1.2.2 The Client ...

Amalgamated Grocers (AmGro) was established 30 years ago. Before then it had been a loose federation of independent grocers, most of whom traded from single locations, but who worked together to achieve the economies of scale available from wholesalers and other suppliers. The economic climate of the time led to them forming a single trading company with 50 stores, mostly in the north of the country. Since then, AmGro has competed successfully with the larger established food retailers and enjoys the largest market share in the north of the England and has the 2nd largest turnover by volume in the country.

They typically adopt a short-term approach to business development, and usually look to “low cost” suppliers. Their record with strategic projects has, unsurprisingly, been poor but they are beginning to recognize that the world is changing radically, and that they need to change their approach. It is worth noting they have not previously had any experience of implementing system for new channels such as the Internet so their expectations may be unrealistic.

1.2.3 Scope of the Home Shopping Order Management System

AmGro believe a significant differentiation in the “home shopping” sector will be the ability to offer a single AmGro-from-Home service over several “channels”:

- **On line Internet.** Customers will be able to browse and order when “on line” and connected to the Internet from their a PC at home or elsewhere. AmGro envisage this will become over time a well-used and popular service, initially favoured by Internet-savvy customers who are happy placing orders of any size over the internet.
- **Telephone Call Centre.** Customers will be able to collect a printed “catalogue” of available goods from AmGro stores, or request one via the post. They will then be able to use these physical catalogues to place orders by telephone through a call centre.
- **Other channels** are being discussed however these are out of scope for now

Whatever the channel, AmGro are absolutely certain that they want a common user experience and appropriate security services in place. Such design factors may have an impact on performance of the system and will need to be taken into account.

1.2.4 Rollout Schedule

AmGro are considering a 5-year period for the take-up and rollout of the AmGro-from-Home system. **In this exercise, you will only be considering the Year 2 Scenario and the performance concerns of the Online Internet channel** (i.e. the Call Centre and other years in the schedule can be ignored).

1.2.5 Input data highlights

Amongst the volumetrics, assumptions and technical definitions which now follows, the relevant Year 2 Online customer figures and data of most interest to your exercise will be highlighted with a **yellow backdrop**.

1.3 Business Volumetrics and Assumptions

1.3.1 Business Volumes Projections

Business & Market Assumptions		Constant	Rollout Years →				
Id	Assumption	Value	Y1	Y2	Y3	Y4	Y5
BA-M-9	Households with Internet access %		31.00%	37.00%	46.00%	49.00%	51.00%
BA-M-10	Internet households with broadband %		0.80%	3.40%	12.20%	24.60%	44.00%
BA-M-11	Maximum AmGro-from-Home households (millions)		1.19	1.46	1.92	2.15	2.39
BA-M-12	AmGro-from-Home uptake as % of total households		10.00%	14.00%	16.00%	18.00%	20.00%
BA-M-13	AmGro-from-Home Online preference %		100.00%	80.00%	70.00%	70.00%	70.00%
BA-M-14	AmGro-from-Home Call Centre preference %		0.00%	20.00%	30.00%	30.00%	30.00%
BA-M-15	Average orders per household per week	1.00					
BA-M-16	Average AmGro-from-Home orders per day		17,054	29,299	43,844	55,391	68,350
BA-M-17	AmGro-from-Home Online orders per day		17,054	23,439	30,691	38,774	47,845
BA-M-18	AmGro-from-Home Call Centre orders per day		0	5,860	13,153	16,617	20,505

1.3.2 Human Users (ARC303)

The “human interfaces” to the AmGro-from-Home (AGfH) system support the access needs of the various users and the various roles they play - they are the **human actors** of the system. This table documents the major human actors for the Order Management system.

User Volumes			Rollout Years →				
Id	User / User item	Description	Y1	Y2	Y3	Y4	Y5
U1	Online user	User using AGfH via Internet (browsing, online order creation and submission, self-registration, ...)					
	Total projected registered AGfH Active Online Users (thousands)		143	197	258	326	402
	AGfH active broadband users (thousands)		1	7	31	80	177
	AGfH active Dial-up users (thousands)		142	190	226	246	225
U2	Call Centre Operators	Call Centre operator handling customer enquiries and order creation via an interactive PC client	0	190	430	540	670

1.3.3 Location Volumes

The following locations have been identified as relevant to the AmGro-from-Home Order Management System. The infrastructure pictures shown in Section 1.4 Technical Architecture Overview show how the infrastructure is to be laid out with respect to these Locations.

Business/System Location Volumes			Rollout Years →				
Id	Location / Location item	Description	Y1	Y2	Y3	Y4	Y5
L1	Store	AmGro Stores (total)	212	228	239	258	265
L2	AGfH Customer Home	Households using AGfH	120,000	210,000	310,000	390,000	480,000
L3	Warehouse	AGfH dedicated dispatch warehouse	6	10	15	19	23
L4	Call Centre	AGfH dedicated call centre		1	1	1	1
L5	Internet Shopping Services	Hosting location for web application	1	1	2	2	2
L6	Central Home Shopping Svcs	Central application location	1	1	2	2	2
L7	Third Party	Third party merchandise location	3	4	5	6	7
L8	Credit Agency	Credit Agency location	1	1	1	1	1
L9	Corporate Services	AmGro Corporate/Legacy systems	1	1	1	1	1

1.3.4 User Usage Assumptions (Year 2)

This table presents a small subset of the assumptions which have been made around user usage patterns for the AmGro-from-Home online service. These have been used in calculating user process volumes.

User Usage Assumptions			
Id	User / User item	Value	Notes
UA-1	Average sessions per order	1.2	Includes logging on to check order status, etc.
	=> Online sessions per day	28,127	
UA-5	Average items per order (online, CC)	25.00	
UA-7	User browsing factor (online)	1.50	Ratio of browses to product selects
UA-8	Orders incorporating shopping list	0.60	Ratio of browses to product selects
UA-9	Shopping list as % of total order	65%	Proportion of items selected from previous/template orders
UA-13	User population growth per year	30%	Gross assumption - could be calculated
UA-15	Peak hour % of total traffic	15%	Assume same for all channels (may be erroneous for Kiosk).

1.3.5 User Process Volumes (Year 2)

A number of core user processes have already been identified in the AmGro-from-Home requirements analysis work. The projected Year 2 volumes of these are shown below, both in term of daily average executions and peak hour executions. It can be assumed that each of these process results in one full web page being returned (which in turn is composed of many elements – see Technical Assumption TA-12).

User Process Volumes

Id	Use Case Scenario	Daily average executions				Peak hour executions			
		Online	Call Centre	Other	Total	Online	Call Centre	Other	Total
UCC1	Logon (customer)	28,127		1,176	29,303	4,219	0	176	4,395
UCC2	Logoff (customer)	28,127		1,176	29,303	4,219	0	176	4,395
UCC3	Logon (CSR)		380		380	0	57	0	57
UCC4	Logoff (CSR)		380		380	0	57	0	57
UCC5	Submit order	23,439	5,860	1,176	30,475	3,516	879	176	4,571
UCC6	Create order	26,043	6,511	1,307	33,861	3,907	977	196	5,079
UCC7	Identify customer		13,283		13,283	0	1,992	0	1,992
UCC8	Retrieve order		6,772		6,772	0	1,016	0	1,016
UCC9	Add shopping list to order	14,063	3,516		17,579	2,110	527	0	2,637
UCC10	Search/Browse for item ¹	307,639	51,273	4,411	363,322	46,146	7,691	662	54,498
UCC11	Add order item	205,092	51,273	2,940	259,306	30,764	7,691	441	38,896
UCC12	Remove order item	58,598	14,649	294	73,541	8,790	2,197	44	11,031
UCC13	Make Delivery Arrangements	26,955	6,739	1,353	35,046	4,043	1,011	203	5,257
UCC14	Make Payment Arrangements	26,955	6,739	1,353	35,046	4,043	1,011	203	5,257
UCC15	Create Shopping List	1,406	664	59	2,129	211	100	9	319
UCC16	Delete Shopping List	1,406	664	59	2,129	211	100	9	319
UCC17	Amend shopping list	1,406	664	59	2,129	211	100	9	319
UCC18	Delete order	2,604	651	131	3,386	391	98	20	508
UCC19	Register new customer	0	0		0	0	0	0	0
UCC20	Delete customer	0	0		0	0	0	0	0
Totals		751,862	170,019	15,492	937,373	112,779	25,503	2,324	140,606

1.4 Technical Architecture Overview

1.4.1 Physical Architecture Overview

The following picture provides a high level overview of the AmGro-from-Home proposed physical architecture (excluding Call Centre infrastructure). For the purposes of this exercise, only the L2_On-line Customer Home and L5_Internet Shopping Services locations are of interest.

¹ Note that this user process also brings back a catalog item large image – see technical assumption TA-14 for the size of this

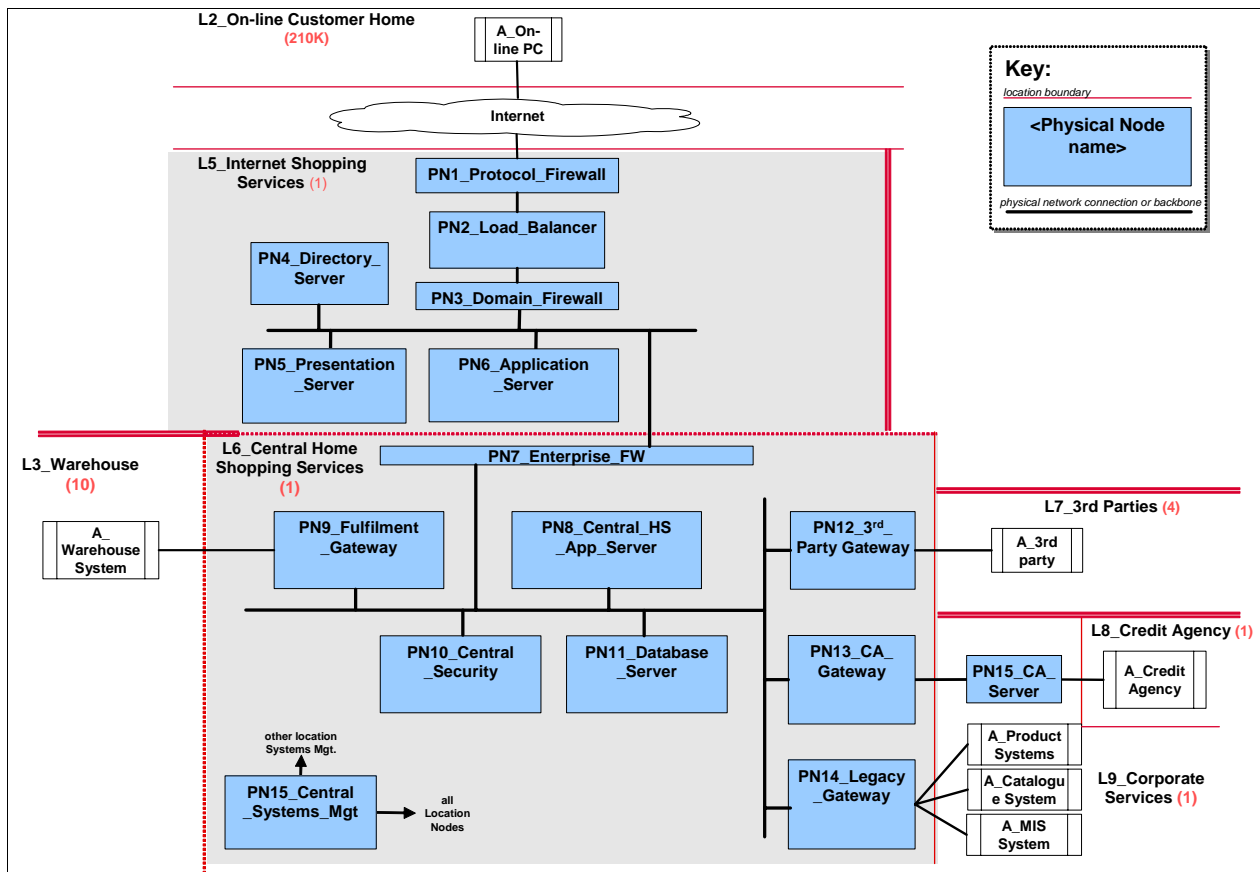


Figure 1: AmGro-from-Home Order Management Physical Architecture Overview

1.4.1.1 Online Channel Locations Detailed View

The following diagram shows a closer view of the L2 and L5 locations. For each node, software and data deployments are indicated – although these details are not important for the purposes of the exercise.

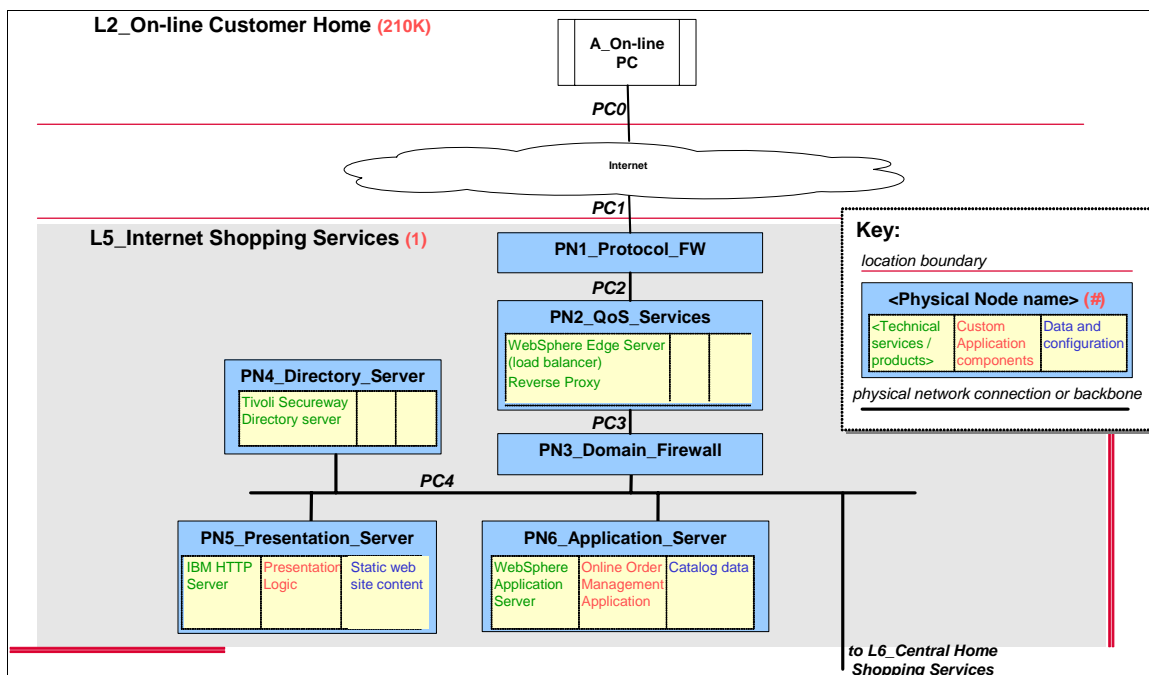


Figure 2: AmGro-from-Home Order Management, Online Channel Locations Physical Architecture

1.4.2 Physical Infrastructure Nodes

The proposed makes, models and performance rating for each of the physical node type in the physical systems architecture are presented in the table below.

Physical Node #	Physical Node Name	Function(s)	Hardware type	Hardware model	# Cpus	Perf. rating	Rating unit	Background load %
PN1	Protocol Firewall	Front protocol firewall	Cisco FW/Router	Cisco PIX 501	1	50	Mbit/s	5%
PN2	Load Balancer	Load balancing and reverse proxy	IBM Intel Server	NetFinity-7000M10-2-400	2	10,391	tpmC	5%
PN3	Domain Firewall	Domain firewall protecting web/application nodes from the DMZ layer	Cisco FW/Router	Cisco PIX 501	1	50	Mbit/s	5%
PN4	Directory Server	Directory and security server	IBM pSeries server	pSeries-p660-6M1-2-500	2	19,692	tpmC	10%
PN5	Presentation Server	HTTP serving	IBM Intel Server	NetFinity-7000M10-2-400	2	10,391	tpmC ²	5%
PN6	Application Server	Web application serving	IBM pSeries server	pSeries-p660-6M1-2-750	2	29,341	tpmC	5%
PN7	Enterprise Firewall	Firewall protecting Central Svcs location	Cisco FW/Router	Cisco PIX 501	1	50	Mbit/s	5%
...								
...								

Note that the required number of each Physical Node type required has not been decided. It may be that more than 1 is required for *throughput* reasons in particular tiers of the architecture – this is for you to decide in the case of PN5.

1.4.3 Physical Connections

Physical connections are the individual networks which are to be used or created to connect the physical nodes in the proposed design. The expected or proposed bandwidth for each connection type is shown in the table.

Physical Connection #	Physical Connection	Relevant location	Main traffic types	Network type	Instances	Capacity	Unit	Background load %
PC0_modem	User Internet Connection (modem)	AGfH Customer Home	HTTP traffic	User's ISP connection	190,194	45 Kbps		2%
PC0_ADSL	User Internet Connection (ADSL)	AGfH Customer Home	HTTP traffic	User's ISP connection	6,694	480 Kbps		2%
PC1	Inbound Internet	Internet Shopping Services	HTTP traffic	ISP connection	1	10 Mbps		5%
PC2	DMZ LAN 1	Internet Shopping Services	Forwarded HTTP traffic	Ethernet	1	100 Mbps		5%
PC3	DMZ LAN 2	Internet Shopping Services	Forwarded HTTP traffic	Ethernet	1	100 Mbps		5%
PC4	Internet Services Location Main LAN	Internet Shopping Services	HTTP traffic; application traffic	Ethernet	1	100 Mbps		10%
...								
...								

² Note: this is a published measure of transactional performance for this server, in tpc-C type transactions per minute.

1.5 Technical Volumes and Assumptions

1.5.1 Technical Transaction Assumptions (Year 2)

In order to calculate the impact of technical transactions on the physical infrastructure, a number of technical transaction assumptions have to be made. This table provides a list of technical assumptions which have been made most relevant to answering the questions – make sure you read them carefully.

Technical Transaction Assumptions

Id	Assumption	Value	Unit
TA-1	Single HTTP element request size	0.5	Kb
TA-2	Single HTTP element response size	2	Kb
...
TA-12	Ave. web page elements per page	30	
...
TA-14	Catalog item large image size	100	Kb
TA-15	Proportion of time a 'Search/Browse for Item' transaction returns an item large image	100%	

1.5.2 Technical Transaction Definitions

This table shows the structure in which technical transactions are defined; as you can see this includes the requesting node, and node of execution for each transaction, and which connection the request and response flow over. Only the technical transaction of interest for the exercise is shown, the "Server HTTP Request" transaction.

Technical Transaction Definitions

Technical Transaction	Source Node #	Source Node Name	Physical Connection #	Physical Connection Name	Execution Node #	Execution Node Name	Cost on execution node
...							
Serve HTTP request	PN3	Domain Firewall	PC4	Internet Services Location Main LAN	PN5	Presentation Server	0.54 tpc-C ³
...							
...							

³ This is the 'cost' or 'weight' of the Serve HTTP request technical transaction in terms of tpc-C like transactions the processor type contained within the PN5 Presentation Server node. This value can be used therefore to estimate the resource which will be consumed by a given number of executions of the transaction on the execution node.



1.5.3 Technical Transaction Volumes (Year 2)

Whereas the above table shows the definition of technical transactions, this table would be used to record the volumes of executions and network bandwidth consumed by each transaction for the target scenario. Again, only the technical transaction of interest, the “Server HTTP Request” transaction, is shown.

Technical Transaction

Totals

On server/network

...

Serve HTTP request

executions in peak hour

total traffic in peak hour (Kbit/s)

??

Presentation Server

??

Internet Services Location
Main LAN

1.6 Response Time Estimation Table

This table is a partially completed end-to-end response time estimation for the user transaction “Search/Browse for Item”. Most of the calculations (including utilisations and response times) for the different components of the system involved in this overall transaction have already been provided.

Hint: In order to answer Question 3:, complete the missing figures and calculations indicated by a red question mark (?) in order to estimate the total end-to-end response time.

Search/Browse for Item – Response Time Estimation

Step	Node	Component	Bandwidth	Msg size (KB)	Utilisation	Service time (s)	Servers (cpus)	Tw (wait time)	Time(s) Processing	Network	Delay
Request											
1	Home PC	Local processing							0.2500		
2	PC0_ADSL/modem	User ISP connection	? Kbps	0.5	0.0	?		?		?	
3	Internet	Internet fixed delay									0.1200
4	PC1 Inbound I/net	AmGro ISP connection	10 Mbps	0.5	0.95	0.0005	1	0.0093		0.0098	
5	PN1 Protocol Fw	Firewall processing	50 Mbps	0.5	0.90	0.0001	1	0.0008	0.0009		
6	PN2 QoS Svcs	Load balance / proxy processing			0.71	0.0012	2	0.0012	0.0023		
7	PN3 Domain Fw	Firewall processing	50 Mbps	0.5	0.90	0.0001	1	0.0008	0.0009		
8	PC4 Internet Svcs LAN	Server LAN	100 Mbps	0.5	0.72	0.0000	1	0.0001		0.0002	
9	PC5 Presentation Server	HTTP Serving			?	?	?	?	?		
Response											
10	PC4 Internet Svcs LAN	Server LAN	100 Mbps	?	0.09	?	1	?		?	
11	PC1 Outbound I/net	AmGro ISP connection	10 Mbps	?	0.95	?	1	?		?	
12	PC0_ADSL/modem	User ISP connection	? Kbps	?	0.0	?		?		?	
Totals:									?	?	0.12
Total response time:									?		



Notes and hints:

- *You can treat network segments as though they were single-server queues for the purpose of the above estimation*
- *The single-server formula for M/M/1 queuing is $T_w = T_s * u / (1 - u)$, where u is the utilisation of the server*
- *A reasonable approximation for a single queue with multiple servers is $T_w = T_s * u^k / (1 - u^k)$. This can be applied to servers with more than one processor, for example.*