3rd Party Proposal: Publication of Near Real Time Data at UK sub-terminals. Modification Reference Number UNC 006

Document Type: Appendices

Ref: 22/06b

Date of Publication: 3 February 2006

Overview:

UNC Modification Proposal 006 was raised by energywatch in November 2004, with the objective of improving the level of transparency in the gas market. In July 2005 the Authority deferred its decision regarding this modification to allow interested parties to consider the incremental value that the release of this additional information would have for the market. This document discusses the further analysis that has been carried out and sets out Ofgem's current conclusions regarding Modification Proposal 006. It also provides interested parties and stakeholders with a further opportunity to comment.

Contact name and details: Hannah Cook, Wholesale Markets

Tel: 020 7901 7444

Email: wholesale.markets@ofgem.gov.uk

Team: Wholesale markets

Deadline for Response: 17 March 2006

Target Audience: This document will be of interest to industrial and commercial gas customers, gas producers, shippers, traders, suppliers and transporters and energy consumer groups.

Ofgem, 9 Millbank, London SW1P 3GE
www.ofgem.gov.uk

Context

In May 2005 Ofgem published an Impact Assessment (IA) on Uniform Network Code (UNC) Modification Proposal 006. In July 2005 the Authority decided to defer its decision to allow the proposal to be assessed against the new baseline following the release of more information to the gas market under the DTI information initiative, agreed with offshore gas producers and National Grid Gas. This IA assesses the proposal against the new baseline of information.

Associated Documents

- Draft Modification Report - 3rd Party Proposal: Publication of Near Real Time Data at UK sub-terminals - Modification Reference Number 0727 - Version 3.0 - 9 February 2005 (and responses)
- Modification Report - 3rd Party Proposal: Publication of Near Real Time Data at UK sub-terminals - Modification Reference Number 0727 - Version 2.0 - 5 April 2005
- 3rd Party Proposal: Publication of Near Real Time Data at UK sub-terminals - Modification Reference Number UNC 006 (0727) - Impact assessment - May 2005 (and responses) http://www.ofgem.gov.uk/temp/ofgem/cache/cmsattach/11579_14305.pdf
- Letter from Steve Smith - Uniform Network Code (UNC) Modification Proposal 006 "3rd Party Proposal: Publication of Near Real Time Data at UK sub-terminals" - 25 July 2005 http://www.ofgem.gov.uk/temp/ofgem/cache/cmsattach/11947_006final.pdf
- Letter from Sonia Brown - Uniform Network Code (UNC) Modification Proposal 006 "3rd Party Proposal: Publication of Near Real Time Data at UK sub-terminals" - 24 October 2005 (and responses) http://www.ofgem.gov.uk/temp/ofgem/cache/cmsattach/13735_October_Letter.pdf?wtfrom=/ofgem/work/index.jsp&section=/areasofwork/wholesalemarketmonitoring
- Letter from Hannah Cook - Publication of Near Real Time Data at UK Sub-Terminals (UNC Modification Proposal 006) - Ofgem Impact Assessment - Case Study - 10 January 2006 http://www.ofgem.gov.uk/temp/ofgem/cache/cmsattach/13767_Jan_letter_merged.pdf?wtfrom=/ofgem/work/index.jsp&section=/areasofwork/wholesalemarketmonitoring

Table of Contents

Appendix 8 – NGG NTS's response to Ofgem's formal request for information regarding costs of implementation

Ms Sonia Brown
Director, Markets
Ofgem
The Office of Gas and Electricity Markets
9 Millbank
London
SW1P 3GE

Mr Richard Court
Manager, Commercial Frameworks

richard.court@uk.ngrid.com
Direct tel
Direct fax +44 (0)1926 656605

www.nationalgrid.com

21 November 2005

Dear Sonia,

Response to Ofgem's Formal Information Request regarding the IS costs associated with the implementation of UNC Modification Proposal 006 (the "Proposal") dated 10th November and Ofgem's letter on the same subject dated 25th October 2005.

Further to your letter dated 10th November 2005, please find the attached document which details the process that was undertaken in order to estimate the IT development and implementation costs for the Proposal and provides the copies of the relevant documentation that you have requested. In addition this letter also provides a response to the letter from Jo Witters to myself dated 25th October 2005 and the subsequent discussions we have had on this subject.

IT Implementation Costs

The Proposal was raised by energywatch and seeks to oblige the publication of "near real time flow data" at sub terminal level.

The cost breakdown previously provided by National Grid NTS shows that the majority of expenditure is predicted to be apportioned almost equally between hardware and software costs. Whilst National Grid NTS acknowledges that some parties have suggested that the estimated IT development and implementation costs are higher than they expected two key drivers for these costs need to be understood.

Firstly, National Grid NTS does not routinely use the two-minute flow data from the sub-terminal meters therefore additional extraction of this data is required from the Integrated Gas Management System (iGMS). IGMS is a core SCADA control system and is an essential system that National Grid NTS uses in managing its primary obligations i.e. to operate the NTS in a safe, secure and efficient manner. The iGMS

system was not designed to supply real time data to non operational systems. All modifications to this (and other core systems) are subject to a planned and controlled release programme that ensures extensive testing and minimal risk of disruption to service and security. Generally these release programmes are planned 18 months in advance of implementation and are limited to no more than one release per year. Essentially a tool is required which can scan and extract the relevant data at the required frequency, whilst ensuring no effect on the performance of the other systems. The cost of the iGMS modifications contribute the majority of the £495k estimated system development and testing costs. Prior to the intent behind the meaning of "near real time" being clarified by the Proposer, National Grid NTS had intended to use hourly data from iGMS, which is more routinely produced and therefore much lower cost to obtain as no new extraction software would be required. Our original estimate of £650k was based on this assumption.

Secondly, analysis has identified that existing web servers used by National Grid NTS are either unsuitable or lacking in available capacity to process and publish the vastly increased volume of data, speed of uplink and the anticipated potential hit rates, even prior to considering acceptable levels of resilience. We have therefore made provision for an entirely new web hardware platform. In addition, the proposed new infrastructure includes provision for a high resilience solution including dual servers, with auto fail-over devices to ensure continuity of service in the event of primary server failure. It is possible for a lower resilience solution to be provided at less cost, but given the current issues raised with regard to the resilience and reliability of our existing web based information provision systems and the fact that we are currently improving resilience on these systems, we believed that the high resilience solution was most likely to be required by the industry in this case. The cost of the new web platform makes up the majority of the £632,500 estimated hardware costs. Lower resilience solutions may cost up to 35% less than this depending upon the specification however this would need to be signed onto by the industry.

The breakdown of these costs, previously provided by National Grid NTS, is as follows;

Cost Element	Cost (£)	Comments
NG Project Resource	153,000	For lifetime of project, including project initiation costs
System Development & Test	495,000	Configure data requirements Extract required data Dispatch data appropriately to Internet Format (and group) data for publication Publish data on web page
Hardware & Configuration	632,500	2 x Load Balancers 2 x Web Servers 2 x Application Servers 2 x D/B Servers 1 Staging Area D/B
Support after	100,000	Application and Infrastructure

Implementation		Support
Total	1,380,500	

Ofgem has previously compared the costs associated with this Proposal with those monies spent or committed on other web projects such as the DTI Information initiative, the ongoing resilience project, and the recently launched Daily Summary Report page. These projects are not directly equivalent to the energywatch Proposal and hence comparisons need to be put in context;

- The DTI Information Initiative, phase 3, published some hourly (and other) information on the existing infrastructure. There were no additional infrastructure costs, all costs were for application development, and additional ongoing support costs to enable 24/7 support. The costs associated with DTI Phase 3 total £162k.

- The resilience project will improve the existing infrastructure resolving many of the performance issues experienced (particularly post DTI) due to the increased hit rate. This will deliver improved reliability and resilience on a phased basis up to June 2006. The cost of this work is estimated at £250,000.

- The Ofgem led DSWG initiative from which the Daily Summary Report was devised makes existing information more accessible to the market by providing a high level summary of the key operational information that was, predominantly, already available from the various National Grid websites with links to the more detailed information, and further provides a new Gas Balancing Alert system. This has been achieved with little hardware expenditure and no new requirement to extract data from our core SCADA systems. The cost of this has therefore been kept to a relatively low £156,000.

The energywatch Proposal is, as explained above, a wholly new type of information provision and will firstly, need to be extracted from core control systems and secondly require new infrastructure because of the increased frequency and volume of information that will be published. Therefore there are differences in the IS requirements between DTI and mod 006. Hence although both may deliver similar information via the internet, the volume and frequency of the mod 006 data requires more IS investment in data extraction and web architecture.

Treatment of Costs in relation to the Proposal

Since Ofgem's letter in July 2005 which deferred any decision on the Proposal until April 2006, and the subsequent letter from Simon Cocks to Steve Smith dated August 2005, National Grid has been considering ways in which it may remain possible to implement the Proposal in line with Ofgem's expectations, i.e. by October 2006.

Following the normal route, of not committing significant expenditure until positive decisions have been received would mean that the data extraction requirement would be too late for the 2006 iGMS release and implementation would therefore not be possible until October 2007. We are currently proactively assessing whether a more generic extraction tool, which includes the required functionality for the Proposal, might also have some broader uses particularly given the increasing

requirements for additional data provision to the industry. If it is decided that this is the case then we could include the extraction tool in the October 2006 iGMS release and remove this long lead-time item from the implementation programme. There still would be risks to the feasibility of meeting the October 2006 implementation date as the web platform is also likely to have a lead time in excess of 6 months. Should we proceed with the generic data extraction tool, this would leave the question of treatment of the iGMS data extraction costs in relation to the Proposal, i.e. should a portion of the costs for the extraction tool be included within the cost / benefit analysis that Ofgem are planning to undertake?

In response to requests from Ofgem and the industry National Grid NTS has investigated the feasibility of obtaining firm IT implementation costs for this project. This would not normally be forthcoming without significant expenditure to create detailed design specifications, etc. National Grid has however recently instructed its service providers to revisit their quotations and provide a greater certainty and explanation to costs. This exercise is anticipated to conclude before the end of next month i.e. in time to inform Ofgem's proposed RIA in January 2006.

Funding for IS Developments

IS development costs are treated as part of National Grid NTS's System Operator internal costs. Therefore not only do we have an incentive to minimise any such costs but it is also in the interests of Shippers that costs are kept to a minimum. National Grid NTS received no explicit allowance for these developments at its last Price Control Review and hence believes that all expenditure is incremental. This being the case, if the generic extraction tool was justified as a broader development then whether this was assigned to the costs of mod proposal 006 or not would have no different effect on the costs that shippers and hence end consumers would face.

Role of xoserve

At the meeting on 10th November, Ofgem expressed surprise that xoserve were not directly involved with the implementation of this proposal. We explained at the time that iGMS system is a System Operator system and therefore was agreed at hivedown to sit within National Grid NTS and not xoserve. We note however Ofgem's view that the provision of data via the internet may fit better with xoserve's role and are happy to discuss ways in which this might be facilitated with xoserve in the future.

National Grid NTS Cooperation with regard to the Proposal and the modification process

A further point raised by Ofgem during our meeting was the extent to which the energywatch proposal could have been developed more fully at an earlier stage. We advised the Proposer prior to the Proposal being raised that the modification needed further development, expressed our view at the initial Panel meeting that the Proposal should not be sent directly to consultation and sought clarification from the Proposer at the development Transmission Workstream meetings in Dec 2004 and Jan 2005. Some clarification was received however a number of Workstream attendees, including ourselves, still believed further work was required. However the

recorded majority of the workstream attendees was, against our view, that the proposal did not require further development and that the workstream should recommend to the Panel that the Proposal be sent for consultation. At the following Panel meeting the Panel voted to send the proposal to consultation. It could be argued that National Grid NTS had the opportunity to raise an alternative proposal however we chose not to do this since we did not believe that we could construct a proposal which we believed furthered the relevant objectives.

The full interpretation of the intent behind the Proposal (i.e. the 2 minute data at 12 minute interval requirement) was not made clear until June 2005. This directly resulted in a revised cost estimate that is now the subject of Ofgem's formal request for information. We believe we have both responded where appropriate at the various stages of the consultation process and also facilitated the development of this Proposal through the workstream and in discussions with the Proposer raising a number of considerations with regard to the Proposal. The IT development and implementation costs is obviously only one, albeit important, consideration with the Proposal. Indeed we have provided Ofgem with a large amount of data on other aspects of this Proposal.

As ever should you wish to discuss any of the information provided in this response or more broadly in connection with this proposal then please do not hesitate to contact either myself or Stephen Johnson on 01926 656200. For example, we would welcome a discussion on potential cost allocation as noted earlier in this letter.

Yours sincerely,

Richard Court
Manager – Commercial Frameworks, Transmission

Appendix 9 - Documents accompanying NGG NTS's response to Ofgem's formal request for information regarding costs of implementation

Explanation of IT Development and Implementation Costs Associated with UNC Modification Proposal 006 Raised by energywatch

1. Summary

Following the raising of Modification Proposal 0727 (later renumbered as UNC 006), costs were identified for delivery of the required data. Due to lack of clarity regarding the proposed interpretation of "Near Real Time" and emerging concerns regarding the resilience of existing internet services for data provision, these original requirements (and associated costing) were unreflective of the delivery which subsequently transpired to be required.

Subsequent requirements and associated costs were identified appropriate to the required delivery and where relevant, the supporting documents detailing these are embedded below, within the commentary.

2. Background to original estimate of £650K

Towards the end of 2004 National Grid provided an original estimate of £650K. This high-level impact analysis was provided on the basis that the energywatch requirement for publication of instantaneous flow data at 'near real time' could be accommodated by evolution of the existing Information Exchange (IE) website, and on this basis costs were derived. At this time, National Grid could not obtain clarification of the 'near real time' definition and assumed that a publication time at hourly intervals would be acceptable. The data required by the energywatch Modification Proposal 0727 would be supplied by the iGMS (Integrated Gas Management System), National Grid's primary network management tool for the National Transmission System, and analysis carried out at this stage indicated that both iGMS and IE could meet hourly publishing intervals.

Subsequent analysis (detailed below) identified that (modified) IE infrastructure would not be able to support the energywatch requirements. Additionally, when confirmation of the publishing frequency requirements (6x2 minute reads of telemetred data, available to the website as a rolling 12 minute update) was received in June 2005, National Grid were able to identify that the energywatch requirements would place an unacceptable load on the operational performance of the iGMS and in fact could only be achieved with significant re-engineering.

3. Requirements Analysis

The energywatch UNC Mod 006 'Publication of Near Real Time data at UK sub-terminals' sought to place obligations on National Grid NTS to publish real time flow data at:

- All NTS entry points that are owned and operated by National Grid NTS
- All entry points that are capable of accepting gas flows at rates greater than 10mcmd

- All individual sub-terminals that are capable of accepting gas flows at greater than 10mcmd.

Responding to the request from Ofgem for the indicative costs associated with this mod, the first stage is clarification of requirements, since those have direct impact on the required hardware and software solution and therefore indicative costs. Embedded below are the expressed requirements. **Please see Appendix 9.a.**

"Reqs for RFI.doc"

This analysis highlighted that current infrastructure could not support the energywatch requirements, and that a new solution should be designed that could meet requirement for publication at near real time, whilst fulfilling the requirement for high resilience, availability and performance.

4. Proposed solution design

As a result of the requirements analysis a new infrastructure was proposed - see embedded documents for details: **Please see Appendix 9.b.**

"AJ Energy Watch Architecture Diagram

"Logical Architecture for RFI.doc"

The design informed supplier estimates for the proposed new infrastructure. The quotation received detailed an estimate of £571,000, Rough Order of Magnitude (Including a contingency amount as per standard industry practice.) for hardware and configuration attached below: **Confidential Information - not included.**

In tandem a respective proposed software (application) solution was defined to meet the requirements and associated costs were identified (Again including a contingency amount as per standard industry practice.), see attached document: **Confidential Information - not included.**

5. National Grid resources

National Grid costs associated with development and implementation were identified and detailed as per the attached spreadsheet. National Grid resources comprise business and IS costs and were derived based on experience of implementing solutions similar in scope and complexity: **Please see Appendix 9.c.**

"Energywatch costs.xls"

6. Final response and Project Mandate

Following the work undertaken to provide estimates for hardware, software and resources, production of the summary document "Publication of iGMS instantaneous flow data via Internet (energywatch request)" was completed – see embedded document below: **Please see Appendix 9.d**

"Publication of iGMS instantaneous flow da

Having supplied the estimate to Ofgem, and in readiness for the start of a significant project, a draft Project Mandate was then produced: **Please see Appendix 9.e.**

"Energywatch Project Mandate.doc"

The Project Mandate is the first deliverable of National Grid's project management methodology, Delivery Excellence, used to enable the process of prioritising and scheduling projects by IS and Business management. However this process was suspended upon receipt of Ofgem's letter dated 25th July 2005 notifying National Grid that Ofgem intended to defer any decision on implementing this change until April 2006.

Appendix 9.a

Publication of near real time flow data via Internet (EnergyWatch request)

As a result of the DTI Information Initiative Transco has been tasked with publishing various 'categories' of data to external parties via the Information Exchange (IE) website. Transco has complied with this directive and 4 'categories' of iGMS-sourced data were published via IE from July 2005.

The energywatch requirements are an extension of this data provision to include the 'near real time' publication of individual sub-terminal, onshore field and storage withdrawal flows.

High Level User Requirements

Functional

1. Data will be sent from the iGMS data publication hub every 2 minutes for energywatch data. Future configurable data will be sent from the iGMS hub every 6, 30 60 minutes and daily.

2. Data details to be published comprise, for each data item (metered flow point for energywatch)

- Some elements of standing data (parameters). This data may change but extremely infrequently (upon equipment modification).
- Some elements of context data (e.g. associated site name)
- Data item value
- Timestamp that data was written to source database
- Timestamp of data acquisition (for publication)
- Timestamp of data publication
- (Where appropriate) A flag/indication that the data is dodgy (this may require some interpretation of data attributes such as flags and timestamp).

3. For energywatch there are currently (not likely to change) 25 flow data items (and associated attributes and context data), that will be received from the iGMS data publication hub every 2 minutes, for onward publication as specified (see 4th bullet) i.e. 18,000 data item values + associated attributes and context data to be published during a gasday. However, the iGMS data publication hub is being designed to cater for the straightforward configuration of further data items, which will be already stored in the data hub, for 'fast response' publication. If all data items held in the hub were configured for publication this would result in the requirement to publish 1,000,000 data items + associated attributes and context data during a gasday.

4. Energywatch have defined their publication frequency requirement as 6, 2 minute reads of telemetered data, available to the website as a rolling 12 minute update i.e. at each 12 minute bar.

5. For energywatch the end-to-end latency requirement, from acquisition in the iGMS Network Manager, to publication on the website is 12 minutes. But, to synchronize with the designed flexibility of the data publication hub component a maximum end-to-end (source system to website) latency of 2 minutes should be targeted.

6. Data to be published (for energywatch) to an appropriate National Grid web page for access by any internet User.

7. There is no current requirement for specific publications to specific Users/recipients.

8. The requirement is to publish in the public domain (see 6th bullet), but for the energywatch data it is anticipated that consumer groups, large End Users and Shippers with no upstream affiliates will have the greatest interest in the data. For other data which will be available in the iGMS data hub, and so can be readily published in the future, the audience may be larger. So for guidance only, the current User metrics apply:

- Total Business Associates (Gemini Users) 192
- Total NTS VLDMCs (large End Users – Transmission) 62
- Total Sub-LDZ VLDMCs (large End Users – Distribution) 56

9. It is not known what hit rates will result from either the provision of the energywatch data, or other data that can be made available from the iGMS data hub. The attached statistics show the use of the current APF report on IE3. This report shows the same data that will be published for energywatch, but in geographically aggregated and hourly form. **Please see Appendix 9.a.i.**

"Hit rates for APF on IE3.xls"

10. Data to be published 24/7

11. Data to be published in formats that allow User to both view (doesn't have to be PDF) and download data

12. The data to be published as soon as it's available (at or beyond the expected time)

- Including late data
- Where no data available from source system, don't publish
- Publication timeliness requirement to be monitored and the Business (including App Support) notified of breach

13. Data published to be available online for 2 years.

Non-Functional

14. Flexibility to choose/configure what data to acquire from iGMS Network Manager and specified Business Apps. (In iGMS Data Publication Hub scope).

15. Flexibility to choose/configure what data to publish

16. Flexibility to choose/configure frequency of publication

17. Flexibility to choose/configure recipients

18. The iGMS performance must not be degraded by the solution for this requirement. (In iGMS Data Publication Hub scope).

19. Future changes in data publication should not incur major performance and regression testing overheads

20. There must be no degradation of performance in the publication of existing IE reports. Transco have financial liabilities associated with the publication of some existing IE reports.

21. If they approve the UNC mod in April, Ofgem require solution implementation on 1st October 2006.

Appendix 9.a.i

Date	View count	Download count	Hit count
01-Sep-05	27	4643	4670
02-Sep-05	7	4354	4361
04-Sep-05	4	4394	4398
05-Sep-05	26	4599	4625
06-Sep-05	7	2465	2472
07-Sep-05	16	5597	5613
08-Sep-05	10	5694	5704
09-Sep-05	9	5479	5488
10-Sep-05	0	4972	4972
11-Sep-05	2	6176	6178
12-Sep-05	10	5350	5360
13-Sep-05	8	5398	5406
14-Sep-05	4	2594	2598
15-Sep-05	36	7037	7073
16-Sep-05	17	5296	5313
17-Sep-05	4	4768	4772
18-Sep-05	11	5563	5574
19-Sep-05	20	5123	5143
20-Sep-05	20	5728	5748
21-Sep-05	15	5372	5387
22-Sep-05	19	5575	5594
23-Sep-05	24	5180	5204
24-Sep-05	11	9966	9977
25-Sep-05	5	6238	6243
26-Sep-05	25	5514	5539
27-Sep-05	73	7322	7395
28-Sep-05	47	5637	5684
29-Sep-05	12	4403	4415
12-Oct-05	0	1059	1059
13-Oct-05	0	105	105
14-Oct-05	5	1218	1223
15-Oct-05	27	634	661
16-Oct-05	28	530	558
17-Oct-05	37	2600	2637
18-Oct-05	19	2762	2781
19-Oct-05	22	2831	2853
20-Oct-05	14	2899	2913
21-Oct-05	10	2828	2838
22-Oct-05	25	2957	2982
23-Oct-05	18	5079	5097
24-Oct-05	38	3152	3190
25-Oct-05	38	2821	2859
26-Oct-05	12	2560	2572
27-Oct-05	14	6508	6522
28-Oct-05	6	7557	7563
30-Oct-05	25	23517	23542
31-Oct-05	51	9673	9724
01-Nov-05	22	9495	9517
02-Nov-05	29	7696	7725
03-Nov-05	31	7131	7162

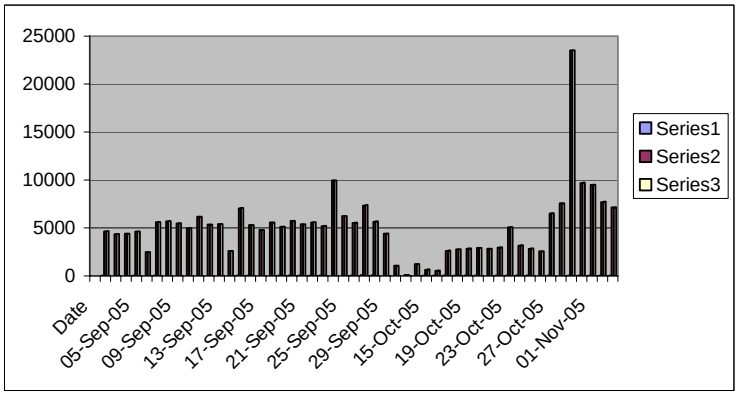

Appendix 9.b

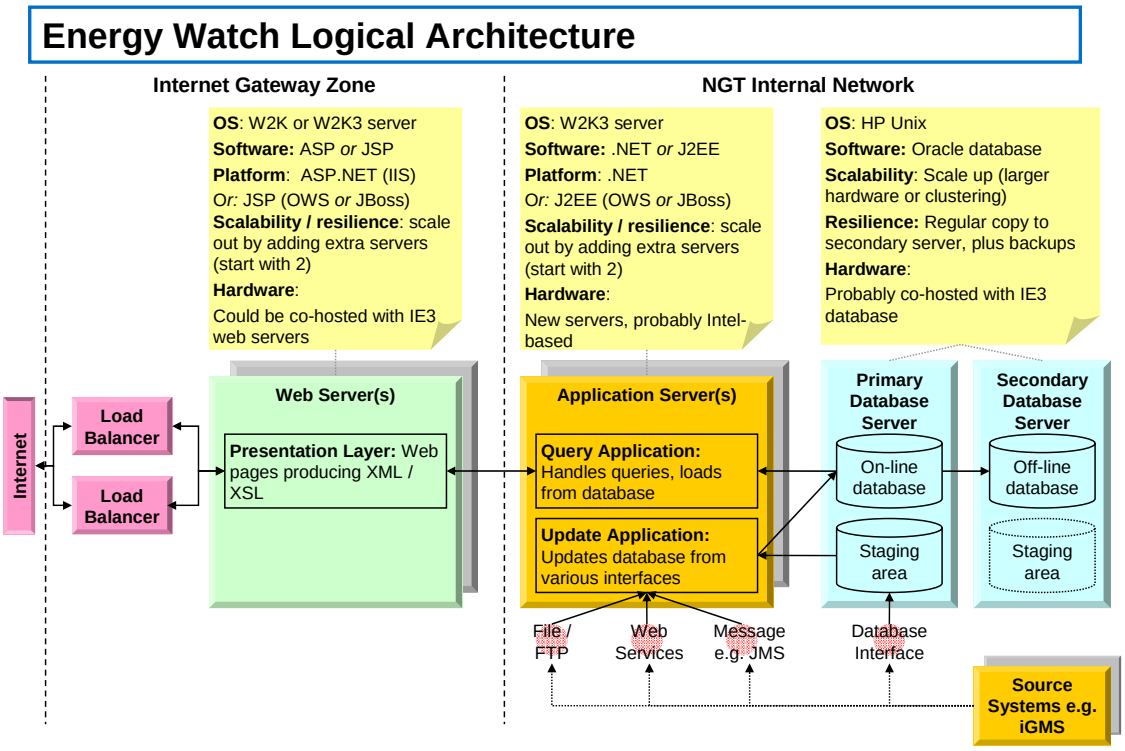

Energy Watch Logical Architecture

1. Introduction

As a result of the DTI Information Initiative NGT has been tasked with publishing information regarding the gas transmission system via a public website. Most of this data will be sourced from iGMS, although other sources might be involved in the future. Several “categories” of data are already published via the Information Exchange (IE3) website, but there is now a requirement to extend this to include “near real time” publication of sub-terminal, onshore field and storage withdrawal flows.

The existing IE3 architecture may not be appropriate to meet the new requirements. This document explores the key requirements which impact on a solution architecture, sets out an appropriate logical architecture (with a number of options), and explores the possible relationship to IE3.

2. Key (Architecturally Significant) Requirements

A number of requirements critically dictate the system behaviour and the architecture(s) which can therefore be used. These are outlined below with current

known values or working assumptions. Each of these must be confirmed, and the proposed architecture reviewed accordingly.

2.1 Data Rates and Volumes

The system will publish data from iGMS at five minute intervals. For each interval the data will be a snapshot of flow rates at each of around 40 metered points. The exact data content is TBC, but is estimates to be 20-100 bytes per reading, a total of 7-48 Kbytes/hour.

The system will retain up to two year's data online, a total of around 8.4 million readings, or less than 1GByte data volume.

2.2 User Numbers, Load and Performance

There are around 100 "expected" users from Shippers and other regular market participants. Public, casual and incidental use is unlikely to more than double this number.

Response times are TBC, but as a base assumption common queries should return within 15s, which is commonly regarded as a standard figure for internet applications.

The system is likely to be "polled" by users, rather than providing an event-driven notification system. Some of this polling is likely to be automated. Ideally each user would poll once in every update cycle, generating a low transaction rate of around 0.7 queries per second.

However, in practice automated queries will run more quickly than this, particularly if certain users believe that they may gain competitive advantage from quickly accessing the latest information. It would therefore be reasonable to design the system to serve one query to each user within the standard response time. This is a rate of 13 queries per second.

It is, of course, possible for users to issue queries in excess of this level. However, NGT should aim manage the load rather than accepting any arbitrary and artificially high demand. A possible approach to this is described under "*3.2 Performance Management*" below.

It is sensible to allow some headroom for application overheads, additional data or query types, and update transactions. Therefore *the system should be designed for 20 transactions (including queries) per second.*

There is no plan to undertake significant post-processing on the data after extraction from the source systems.

The maximum permitted latency from source system update to publication is five minutes.

Use of this system should not impact the performance of the iGMS or the existing IE reports.

2.3 Availability and Reliability

The system should be "continuously available", with no protracted planned outages. It is not an operationally critical system, but regulatory implications of any significant unplanned outage are unknown. However, given the five minute data change cycle an interruption of a 1-2 times this duration should be readily tolerable (TBC).
This is predominantly a query-only system with few or no user updates to data. Therefore there is no need to persist user session information through any failure, and simple "batch" transaction techniques should be adequate to protect update processes.

2.4 Security

The system will be a "Public access" system, with all data available to all users. There will not therefore be any requirement to authenticate users or manage identity information beyond that required for simple usage statistics.

2.5 Flexibility

If possible without significant extra cost the system architecture should allow for several expected dimensions of future flexibility:

- Use for publishing other data, possibly from different source systems and delivered via different interfacing mechanisms,
- Improvements to data latency and publication rates. The current assumption is that 1 minute would be a reasonable minimum target for each.
- Improved "speed to market", i.e. the ability to quickly publish similar data without the need for a protracted system development. In particular, it must be possible to change data sources and publication formats without a major regression test exercise.

The data format should allow the users to easily review data on-screen and also to download and further process the same data.

2.6 Delivery timescale

The latest date for delivery is July 06, but Ofgem would prefer something sooner, e.g. Oct 1st 2005.

3. Logical Architecture

3.1 Overview

The following diagram summarises the proposed solution, which is a standard 3-tier web application using modern component technologies:

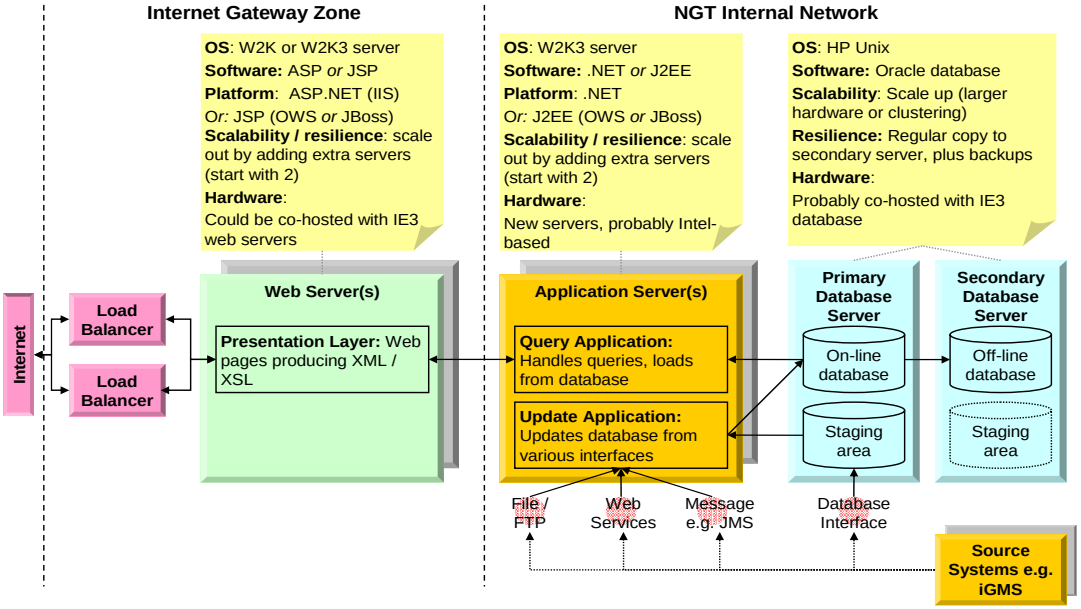

The web server focuses on presenting the web pages to the user, navigation logic, and output formatting. It connects to the application layer in the middle tier, but does not connect directly to the database. The “presentation” application will use a combination of static HTML pages, Active Server Pages (or Java Server Pages), and XML. Query results will be formatted as XML, which can then be either presented to the user using XSL stylesheets, or downloaded and further processed if required. The middle tier will be written using either Enterprise Java (J2EE) or Microsoft .NET component technologies (to match the presentation layer). There will be two “applications”:

- The query application. This will receive query requests from the front-end and return responses. Where necessary this will interact with the database to extract appropriate data, but in many cases the response will be constructed from data held in active objects.

- The update application. This will service the incoming interface(s), processing incoming data and submitting updates to the database. A variety of interface mechanisms can be supported within a consistent processing architecture. Oracle-based interfaces may deliver data directly to a “staging” area – the update application will treat this area in the same way as any other incoming interface, processing the data and managing updates to the main on-line database.

Both applications will share a “data access layer” which will manage interactions with the database. That will be a relatively “dumb” Oracle database hosted on a primary server, with a warm-standby secondary server maintained from the primary by a background data mirroring process.

The use of an Oracle database is a sensible “default” option, but not the only one. Data volumes are sufficiently low that other alternatives, including file-based storage, could also be considered.

3.2 Performance Management

The expected transaction rate suggests a multi-tier system, with an intermediate application layer managing active data and database connections independent of the front end. While it is possible that the performance requirement might be met from a single server at each layer, the availability requirements dictate at least two. Therefore the front and middle tiers are designed to use two or more relatively small servers, and throughput can be increased by "scaling out", adding more identical servers at the appropriate layer.

The database server should be lightly loaded in this design, and the databases may sensibly be co-hosted with other Oracle services on appropriate large shared servers. In the unlikely event it is necessary increased capacity can be achieved by scaling up or clustering those servers.

3.2.1 Demand Throttling

NGT must learn from previous experience and manage user demand so that the system is not "thrashed" by unlimited automated polling of the web pages. The requirement to publish the data via simple public web pages prevents a true event-driven approach (e.g. publishing data to message queues). Failing to properly manage demand will destroy any "level playing field". Therefore the application must impose a degree of demand management itself. Various techniques are possible. The following summarises one simple possibility:

- Users access the data via a "gateway" page. This includes the time of last publication, and a link to the full query page. The gateway page is a "static" HTML page (updated once every five minutes) so that it can be polled frequently with impunity.

- The users can only get the current data by waiting for the gateway page to change, then following the link. Re-polling a previous query page will only deliver old data.

- The link contains an "encrypted" element which changes at each publication, so that users cannot "predict" the address of the new data and repeatedly poll it.

3.3 Availability and Reliability

The proposed architecture should require very few planned outages, and any outage should be short, limited only to the time taken to reconfigure the system's connectivity (i.e. which servers are active at each layer) and restart appropriate processes.

The database can be backed up from the off-line copy on the secondary server. There is no persistent data on any other server, so backups need only be taken occasionally, probably with the system in use.

There is no single point of failure. Any component can be taken out of service for maintenance or updates. Most updates can be applied to one node first, (while the others continue to support the live service) then the other(s). The proposed component technologies both also support “hot” updates to the application components if that is appropriate.

3.4 Data Management

The database in this architecture is used to “persist and populate” object data, rather than to serve every query. It also has a low projected total volume. This allows it to have a very simple structure, and to be managed with simple tools.

More complex management software such as Oracle Warehouse Builder should not be required, but could be introduced later if it becomes necessary to manage more complex database interfaces.

There is probably no case for a complex, “heavyweight” data access mechanism serving the application layer. To reduce complexity the detailed design should consider adopting a lightweight persistence framework (such as Hibernate for a Java-based design).

3.4.1 Non-database option

The projected data volumes and data update activity are relatively low. The database should be co-located with others on suitable shared servers. If there are any barriers to this approach it may be appropriate to consider alternatives such as an XML file-based store. Adoption of an abstract persistence framework such as Hibernate would allow the data store to be changed independent of the rest of the application, which might also be useful for development and test purposes.

3.5 Interfacing

The proposed architecture would support a variety of interface mechanisms, via interface-specific handlers “plugged in” to a common set of components preparing and handling updates to the database.

The best mechanisms would be file-based (including FTP), message-based (e.g. JMS) or web-service (SOAP) based, allowing the source system to be loosely coupled with the Energy Watch system, and allowing updates to queue up for later processing in the event of a system or network outage or slowdown. JMS is probably the best option given the Unix/Java basis of key source systems.

The architecture also supports interfaces which deliver data directly to a “staging area” on the Oracle database. This may appear easier where the source system also has an Oracle database, but has a number of implications:

- The source system must either produce a stream of updates (e.g. via database update triggers) or allow regular queries. Either may have some impact on the source system.

- Direct database updates require both systems to be available concurrently, which may not be the case where the source system has a very high availability requirement (e.g. iGMS). The option is to use Oracle-based queuing (AQs), but this increases complexity.

Discussions about how best to serve the data from iGMS are ongoing and a final recommendation will follow.

3.6 Platform Choice

This architecture could be built readily using either J2EE or Microsoft technologies. The platform choice will depend largely on supplier skills, and the proposed implementation approach.

The only significant technical indicator for J2EE is if iGMS uses it and can easily drive JMS, although it would be possible to bridge an MS-based system to JMS if required. The other interface options (database, file and web services) are all platform-independent. Performance might also be a discriminator (small MS-based systems tend to out-perform J2EE) but this is unlikely to be critical. A final decision is not needed until issues of approach and sourcing have been resolved.

4. Relationship to IE3

IE3 is designed to deliver complex Business Objects reports on a relatively long cycle (currently every 10 minutes). The core application servers are quite heavily loaded, and there are already performance issues. The application architecture is not suited to delivering smaller data updates on a short cycle - the possibility of changing the IE3 update cycle to 5 minutes or shorter has been reviewed by a third party contractor which did not think this could be easily achieved around the software versions currently installed.

The Energy Watch application cannot be delivered “under IE3”, as the application architectures are quite different. However, there are some opportunities for infrastructure rationalisation. The physical architecture is strictly not an NGT issue but it is possible to propose an obvious solution which will minimise any new infrastructure:

- The Energy Watch presentation layer could be co-hosted on the same web servers as IE3, sharing the same load balancing and (if required) user management services (further analysis required).

- The database could be co-hosted on the same database servers (further analysis required).

- The application layer could be hosted on new servers dedicated to Energy Watch. For the production system these will be two small-medium servers, probably Intel-based.

- A similar approach can be taken for test environments.

Appendix 9.c

EnergyWatch				
Cost	Estimate	Risk for estimate	Total	Scheme Values
Scheme initiation costs	£30,000	20%	£36,000	£36,000
NGT IS & Business Resource Costs	£97,000	20%	£116,400	£117,000
Supplier quotation for development costs	£450,000	10%	£495,000	£495,000
Hardware & configuration	£575,000	10%	£632,500	£632,500
Run the business (Annual opex)	£100,000	0%	£100,000	£100,000
		Total	£1,343,900	£1,380,500

Notes
Annual Opex = 1 Application Support + 1 Infrastructure Support @ £35,000/year (from Lesley Wood IE3 resilience project)

Appendix 9.d

Publication of iGMS instantaneous flow data via Internet (EnergyWatch request)

System Development

The provision and publication of telemetered instantaneous flow data at Transco system entry points at near real-time frequency (the EnergyWatch requirement) impacts two Transco systems, the SCADA component of the Integrated Gas Management System (iGMS) which provides the data, and the Information Exchange System which currently supports the publication of all data and reports to Transco's Information Exchange website.

Considerations

Whilst determining a proposed solution for the EnergyWatch requirement consideration has been given to the following points:

- *The performance and availability of the iGMS must not be compromised.* The iGMS SCADA component currently provides safety critical data directly to Control Room staff, and business critical data and information via the iGMS business applications to Users throughout OT&C. The current data provision requirement, which is supported by the iGMS design, is for (relatively) low frequency report production. With consideration for this new high frequency EnergyWatch requirement the proposed solution distances the new data provision function from existing iGMS functionality as much as possible whilst protecting the resilience of this new data flow.

- *The future iGMS release schedule should not be compromised.* The development and release of the proposed EnergyWatch solution must be managed alongside future planned enhancements to iGMS, which include safety and regulatory changes such as supporting functionality for both the Transco Emergency Strategy Planning tool and the new Exit Capacity Regime.

- *The performance of the existing Information Exchange system should not be compromised.* Transco has existing obligations to publish data at predefined frequencies to the Information Exchange website. The Information Exchange system was designed to cater for the existing requirement of low frequency report publication. For these reasons it has been assumed that for the EnergyWatch requirement a publication system would be built alongside, but separate from, the existing Information Exchange functionality to provide a higher frequency publication, highly resilient, highly available solution with minimized risk of data loss.

- *Equal access for all potential Users.* The data presentation has been designed to maintain equal access to all Users rather than larger Users being able to

repeatedly query the data and in doing so effectively block the access of less sophisticated Users.

Scope

- The proposed system is a simple data relay and publication system. The data acquired from the SCADA system at 2 minute intervals is batched and made available for publication every 12 minutes, with a latency of 12 minutes before being visible on the website e.g. data acquired between 11:00 and 11:12 Will be visible on the website from between 11:23 and 11:24. This meets the EnergyWatch requirement as received by Transco on 16/6/05.
- The data to be published is the instantaneous flow data for each Transco system entry point that is capable of fulfilling system entry flow criteria. Note that these criteria have not yet been defined and communicated to Transco by EnergyWatch. The number of eligible entry points is assumed to be < 40.
- The data to be published will comprise, for each eligible point, the instantaneous flow value, timestamps and a simple data status indicator.
- All data published will be available, in its original published form, for User access for 2 years beyond its original publication date.
- The data will be published for public access only.
- Users will be able to view and download the published data.
- It is assumed that a maximum of 100 Users will view or download the data concurrently.
- The solution proposed here is high resilience with auto-restart failover capability.
- A half hour planned outage will be required daily for housekeeping. Other planned outages will be confined to major upgrades of infrastructure or software.
- Failure to publish data will occur only in the event of source system failure i.e. iGMS or telemetry system.

Costs

The costs quoted below are for the build, delivery and ongoing costs of an independent high resilience, auto-failover solution, with new hardware throughout. The costs also assume delivery separate from existing IS release schedules (reference *Considerations* above). It may be possible to reduce costs e.g. using existing commissioned hardware to host some aspects of the proposed solution, or by amalgamation with existing planned IS releases. However, significant work must be undertaken to understand the risks and constraints resulting from any co-hosting (reference *Considerations* above).

Element	£Cost
1. Scheme Initiation	36,000
2. Transco Resources	117,000
3. Software Development	495,000
4. Hardware and Configuration	632,500
5. Annual Running Costs	100,000

Potential Timescales

It is expected that the proposed solution described above could be delivered in Q4 2006. This timescale is with reference to existing iGMS work being undertaken, as described in *Considerations* above.

Appendix 9.e

PROJECT MANDATE		Document Ref & Version No:
Programme:	Project: Telemetered data provision and publication to Transco NTS website for energywatch	
Author: Angela Page/Steve Pownall		Date: 29/6/05

Purpose:

This document has been produced to capture an initial business view of the objectives, benefits and indicative costs for this project concept, so that the project concept can be prioritised, and potentially investigated further. The Project Mandate will provide the basis for the Delivery Unit, in consultation with the relevant business unit, to decide to authorise further development of the project concept into a Project Brief during the Starting Up a Project Stage.

Project Definition:

What type of project is this?	Full Project
Is there an infrastructure element to this project?	Yes
Project Name	Telemetered data provision and publication to Transco NTS website for energywatch
IS Delivery Unit / Project Area	Operations, Trading & Commercial
Project Executive	Richard Court
Business Analyst	Angela Page
Project Investment Driver	Regulatory
Utility Affected	Gas
Type of Work	Regulated
Project Funding	Business
Chargeable Customer	NTS SO incentive 60/40
Estimated Date of Project Sanction	Subject to Ofgem directive
Estimated Project Start Date	Subject to Ofgem directive
Estimated Project Completion Date	Q4 2006

Background:

energywatch has raised a Modification Proposal (0727) “Publication of Near Real-time data at UK Sub-terminals”. This proposal is to extend the information that is published by Transco NTS as part of the ‘DTI Information Initiative’, specifically the DTI Category 1 report – sub-terminal flow data, currently published on an aggregated, national (North/South) basis. This proposed extension requires Transco NTS to provide and publish the telemetered instantaneous flow data at Transco NTS system entry points at near real-time frequency. energywatch have recently defined

near real-time data as 6, 2 minute reads of telemetered data, available to the website as a rolling 12 minute update i.e. at each 12 minute bar.

This energywatch Modification Proposal has been submitted to Ofgem, who have in turn produced a Regulatory Impact assessment (RIA). This RIA requested a Transco NTS response, to include an impact assessment for IS delivery costs.

Business Solutions have engaged with an solutions architect and hardware and software suppliers to produce a high level proposed systems solution and compile a 'worst scenario' quote for systems costs. This was summarised by Commercial and included in the Transco NTS response to the RIA. This response was sent to Ofgem on 24/6/05.

It is the view of Commercial that Ofgem will direct Transco NTS to implement this modification. It is their view that Ofgem will approve the modification next month (July 2005).

Objectives:

- The objective of the project is to satisfy the energywatch requirement and Ofgem directive to publish telemetered gas system entry flow data to the Transco NTS website at near real-time frequency, and to do so in the most cost-effective way.

- The solution delivered must give flexibility for easy development (short lead time, low cost) to expand, reduce or change the telemetered items in the dataset for publication, or to publish data items from other areas of iGMS.

- The solution has to fit within NGT's strategic approach for data provision to external parties over the internet.

Scope & Interfaces:

The provision and publication of data via the internet will not impact any existing Business process.

In Scope

- Data for publication will be acquired from the iGMS Network Manager component and published to the Transco NTS Information Exchange website. The data is required as 2 minute readings and some context data, and is to be published every 12 minutes as 6 sets of 2 minute data. The publication process delay (latency) has been defined as 12 minutes.

- The existing systems in scope for enhancement are iGMS and Information Exchange. No other existing systems are impacted.

- Because of the increase in frequency of acquisition and publication of data above and beyond what is currently available the existing system framework will need to be extended to include extra infrastructure elements.

- Data published will be available online for 2 years.

- The system availability and support should be 24/7.

The design must be mindful of likely future flexibility required for data provision requirements.

Out of Scope

- Published data cannot be amended on the internet by the User.

- Existing data will not be migrated from iGMS for visibility on the internet.

- Corrections made in the iGMS (Network Manager) to data that has been published will not subsequently be re-published to the internet.

Interfaces to other Projects

- IGMS Releases 3 and 4

- IE3 Resilience (Re-platforming)

- Strategic platform for data provision to external parties

Constraints:

- Delivery timescales are constrained by whatever date Transco NTS agrees with Ofgem.

- The success of the project is dependent on the quality of data being published, which is constrained by the quality of metering and telemetry.

- The current system design (of iGMS or Information Exchange) will not support the enhancements required to deliver this extra functionality.

- The design of the solution will be influenced by the need to not degrade the performance of either the iGMS or Information Exchange systems.

- The solution must be seen to support timely and equitable provision of data (any imposed delay to data access will not be approved by Ofgem).

- The solution delivery date is constrained by the existing iGMS release schedule.

- For operational reasons no iGMS releases are scheduled during the winter months.

Quality Expectations:

Because this work is being undertaken as a result of a directive from the Regulator the delivery timescales agreed with Ofgem will be a critical success factor for this project.

The project will deliver a commercial data provision service via the internet, so the availability of that service to external Users will be the most highly visible measure of success. It will also be key that the data provision is appropriate for ease of use by the end User.

The quality of data published is expected by Ofgem and energywatch to be appropriate to inform markets with a reasonable degree of accuracy.

The project deliverables and process, and delivered solution must adhere to IS internal quality rules (to be advised by IS Stakeholders).

Assumptions:

- The progression of this project assumes that Ofgem approve the energywatch Modification Proposal (0727).
- For costing hardware development elements of the project it has been assumed that all hardware used will be procured specifically for this project. As part of the project initiation phase work must be undertaken to determine whether the existing infrastructure could be used to fit the requirement profile. It is expected that this will be the case but quantitative evidence that co-hosting systems will not be adversely affected.
- For costing purposes it has been assumed that a high resilience solution is required. The quoted costs include inter-site duty/warm standby machines with dual links throughout and auto-failover capability. This assumption has been made because that true availability requirement has not been defined by the Business yet.
- The quoted costs assume a maximum of 100 concurrent Users. If there is a requirement for greater concurrent usage then the quoted costs will need further development.
- It has been assumed that data will be published to all Internet Users on a public basis.
- It is assumed that Transco NTS will wish to facilitate equal access to data by all Users but will design a solution that also includes the consideration of allowing access to the data via End User scheduled scripts.
- It has been assumed that only a very low level of source data interpretation is required. The detail of data to be published has not yet been defined.
- It has been assumed that data will be accessed on a ‘read-only’ basis.
- The cost timetable assumes that delivery in Q4 2006 will be acceptable to Ofgem. This delivery date is itself impacted by the future iGMS release schedule. It is assumed that all initiation work can be completed in 2005 and that development,

testing and any necessary procurement and set-up can be achieved within 34 elapsed weeks.

Outline Business Benefits / Business Case:

The Business benefit resulting from the successful implementation of this project will be the protection and potential enhancement of Transco NTS's reputation. If the energywatch requirement is satisfied to the Ofgem timescales than Transco NTS will not be in breach of its Uniform Network Code obligations. If the requirement is satisfied in a User-friendly, equitable and high availability way then Transco NTS's reputation as a data provider will be enhanced with energywatch, Ofgem and gas market operatives.

Benefits:

- Compliance with Regulatory directive. Required in order to meet UNC obligations.
- Enhanced Transco NTS reputation

Indicative Costs:

Project

Scheme Initiation	£36,000
Transco NTS Resources during Development	£117,000
Software Development	£495,000
Hardware and Configuration	£632,000
Ongoing Costs	
Annual Running Costs	£100,000

Summary:

Financial Projections (ENTER IN MILLIONS OF POUNDS)						
Estimated Five-Year Project COST Profile (£m)						
Financial Year	2005	2006	2007	2008	2009	Total
Opex	0.000	0.000	0.000	0.000	0.000	0.000
Capex	0.451	0.829	0.000	0.000	0.000	1.280
Total	0.451	0.829	0.000	0.000	0.000	1.280
Estimated Five-Year Additional "Run-The-Business" COST Profile (£m)						
Financial Year	2005	2006	2007	2008	2009	Total
Opex	0.000	0.000	0.100	0.100	0.100	0.300
Capex	0.000	0.000	0.000	0.000	0.000	0.000
Total	0.000	0.000	0.100	0.100	0.100	0.300
Estimated Five-Year Project BENEFIT Profile (£m)						
Financial Year	2005	2006	2007	2008	2009	Total
Opex	0.000	0.000	0.000	0.000	0.000	0.000
Capex	0.000	0.000	0.000	0.000	0.000	0.000
Total	0.000	0.000	0.000	0.000	0.000	0.000

Preliminary Risk Assessment:

There is a risk that…	Because of…	Leading to…
Non delivery of project	Lack of resource or funding	Breach of UNC obligations and financial penalties Damage to Transco NTS's reputation
Deliver project with a scheduled iGMS release and iGMS release delayed (and consequently project delivery delayed)	IGMS development delays	Breach of UNC obligations and financial penalties Damage to Transco NTS's reputation
Project impacts iGMS performance or availability	Addition of extra functionality and load on system	Safety and security of Supply compromised
Project impacts Information Exchange performance or availability	Addition of extra functionality and load on system	Breach of UNC obligations and financial penalties Damage to Transco NTS's reputation
Quality of data published is poor	Metering and telemetry issues	Damage to Transco NTS's reputation

Stakeholders:

Business Stakeholder	Contact Name	Interest
Commercial – Commercial Frameworks	Steve Pownall	Project Facilitator
Commercial – Governance & Implementation	Kevin Broadbent	Governance
Operations and Trading	Simon Fairman	Business owner IGMS
Operations and Trading	Paul Gallagher	Business Owner of IE
IS – Delivery	Tony Eyrl	Project Delivery
IS – Technology	Eddy Self	Technology Adopted
IS - Strategy	Derek Walker	Strategic Approach

Document Control:

AMENDMENT HISTORY

Version	Date	Remarks

AUTHOR

Name	Position / Department

DISTRIBUTION

Name	Position	Reason for Distribution
	Business Planning & Governance	To add the project to the IS Order Book as a Planned Project.

DOCUMENT REFERENCES

Title	Document Owner	Document Description	Document Location

DOCUMENT Change History
Owner/Author: Gillian Woodrow
AMENDMENTS

Issue	Date	Change Details
Issue 1.0	21/09/04	
Issue 1.2	3/11/04	Added more prompt text and added formulae back in to the tables.

Next Review Date: 31/03/05
Paper Copies are Uncontrolled

Appendix 10 NGG NTS updated response to Ofgem's formal request for information regarding costs of implementation

Rough Order of Magnitude Costs to Develop Solution for UNC Mod 006 – 16/12/05

Background

Ofgem requested a RIA in June 2005 asking for an estimate of costs associated with making specific operational data available in the public domain at 'near real time'. Transco responded with a quote of £1.4M.

Since then National Grid (previously Transco) has undertaken further analysis and devised an approach for the required work to include:

- The separation of the development into 2 distinct components for a more flexible solution
  - IGMS Data Publication Hub to extract data from O&T Operational System
  - Web Info System to format data and publish to the Internet
- Sanctioned costs for the Data Publication Hub component and started development
- Refined the requirements for the Web Info System to cover the re-use of infrastructure where possible and the scope statements attached. **Confidential Information - not included.**

Summary of Costs for energywatch UNC mod 006

Cost Element	Cost (£)	Comments
IGMS Data Publication Hub	500,000	Sanctioned cost
Web Info System - System Development & Test	429,000	ODC quote
Web Info System – NG IS costs	76,000	
Web Info System - Hardware & Configuration	401,000	Supplier quote
Web Info System - Support after Implementation	43,000	Application and Infrastructure Support
Total	1,449,000	

Detailed Cost Calculations **Please see Appendix 10.a.**

"EnergyWatch costs 20051216.xls"

Detailed Cost Quotes. **Please see Appendix 10.b and 10.c - only certain elements are included due to confidentiality.**

"ROM010 ROM iGMS Data Publication - En

Appendix 10.a

EnergyWatch - Web Info System				
Cost	Estimate	Risk for estimate	Total	Scheme Values
NGT IS & Business Resource Costs	£69,000	10%	£75,900	£76,000
Development Costs	£390,000	10%	£429,000	£429,000
Hardware & configuration	£364,000	10%	£400,400	£401,000
Run the business (Annual opex)	£43,000	0%	£43,000	£43,000
		Total	£948,300	£949,000

Notes

Annual Opex = 1 Application Support @ £35,000/year (from Lesley Wood IE3 resilience project) + 1 Infrastructure Support from Vendors @ £8,000/year (local support covered by existing CSC contract)

EnergyWatch - Data Publication Hub				
Cost				Scheme Value
iGMS Data Hub component				£500,000

EnergyWatch - Total Costs				
Cost				Scheme Value
Total Scheme				£1,449,000

EnergyWatch Response - Resource costs							
Team	Phase	Activity	Grade	Effort (days)	Day Rate	Cost	Sub Totals
NGT IS	Initiation	Project Manager	TM A	15	£279	£4,185	
NGT IS	Initiation	Business Analyst	TS 6	20	£245	£4,900	
NGT O&T	Initiation	Business User	TS 6	10	£245	£2,450	
NGT O&T	Initiation	Solutions Architect	TS 6	10	£245	£2,450	
							£14,000
NGT IS	Delivery	Project Manager (½ for duration and a bit)	TM A	85	£279	£23,715	
NGT IS	Delivery	Business Analyst (½ for duration)	TS 6	70	£245	£17,150	
NGT IS	Delivery	Business User (½ for ½ duration)	TS 6	35	£245	£8,575	
NGT IS	Delivery	Solutions Architect	TS 6	20	£245	£4,900	
							£55,000
					Total	£54,340	£69,000

Questions
What about PMO effort etc.?
Do we need more BAs?

No costs included for development environments

Grade	Hour rate	Day rate
TM A	£37.74	£279.28
TM B	£61.74	£456.88
TS 6	£33.20	£245.68
TS 5	£28.65	£212.01
TS 4	£20.43	£151.18
TS 2	£14.32	£105.97

Appendix 10.b

Publication of near real time flow data via Internet (EnergyWatch request)

As a result of the DTI Information Initiative Transco has been tasked with publishing various 'categories' of data to external parties via the Information Exchange (IE) website. Transco has complied with this directive and 4 'categories' of iGMS-sourced data were published via IE from July 2005.

The energywatch requirements are an extension of this data provision to include the 'near real time' publication of individual sub-terminal, onshore field and storage withdrawal flows via IE.

The requirement comprises 2 main components. The first component is the iGMS Data Publication Hub, and this has already been specified. The high level requirements for the second component – to receive data from the hub and publish onwards to the internet – are stated below.

It is important to note that the iGMS data publication hub has been specified to hold energywatch data and extra pre-defined data from various iGMS components. The intention was to build a solution for energywatch and also the capability to respond quickly and cheaply to future data provision requirements. The publication component described below must match this capability for the Business to realise the benefit of the designed flexibility and capability of the hub. Therefore the following requirements cover both the specific energywatch requirements and the necessary flexibility capability.

High Level User Requirements

Functional

1. Data will be sent from the iGMS data publication hub every 2 minutes for energywatch data. Other configurable data will be sent from the iGMS hub every 6, 30 60 minutes and daily.

2. Data details to be published comprise, for each data item (metered flow point for energywatch)

- Some elements of standing data (parameters). This data may change but extremely infrequently (upon equipment modification).
- Some elements of context data (e.g. associated site name)
- Data item value
- Timestamp that data was written to source database
- Timestamp of data acquisition (for publication)

- Timestamp of data publication

- (Where appropriate) A flag/indication that the data is dodgy (this may require some interpretation of data attributes such as flags and timestamp).

3. Some simple data interpretation will be required e.g. unit conversion

4. For energywatch there are currently (not likely to change) 25 flow data items (and associated attributes and context data), that will be received from the iGMS data publication hub every 2 minutes, for onward publication as specified (see 4th bullet) i.e. 18,000 data item values + associated attributes and context data to be published during a gasday. However, the iGMS data publication hub is being designed to cater for the straightforward configuration of further data items, which will be already stored in the data hub, for 'fast response' publication. If all data items held in the hub were configured for publication this would result in the requirement to publish 1,000,000 data items + associated attributes and context data during a gasday.

5. Energywatch have defined their publication frequency requirement as 6, 2 minute reads of telemetered data, available to the website as a rolling 12 minute update i.e. at each 12 minute bar.

6. For energywatch the end-to-end latency requirement, from acquisition in the iGMS Network Manager, to publication on the website is 12 minutes. But, to synchronize with the designed flexibility of the data publication hub component a maximum end-to-end (source system to website) latency of 2 minutes should be targeted.

7. Data to be published (for energywatch) to an appropriate National Grid web page for access by any internet User.

8. There is no current requirement for specific publications to specific Users/recipients.

9. The requirement is to publish in the public domain (see 6th bullet), but for the energywatch data it is anticipated that consumer groups, large End Users and Shippers with no upstream affiliates will have the greatest interest in the data. For other data which will be available in the iGMS data hub, and so can be readily published in the future, the audience may be larger. So for guidance only, the current User metrics apply:

- Total Business Associates (Gemini Users) 192

- Total NTS VLDMCs (large End Users – Transmission) 62

- Total Sub-LDZ VLDMCs (large End Users – Distribution) 56

10. It is not known what hit rates will result from either the provision of the energywatch data, or other data that can be made available from the iGMS data hub. The attached statistics show the use of the current APF report on IE3. This report

shows the same data that will be published for energywatch, but in geographically aggregated and hourly form. **Please see Appendix 10.b.i.**

"Hit rates for APF on IE3.xls"

11. Data to be published 24/7

12. Data to be published in formats that allow User to both view (doesn't have to be PDF) and download data

13. The data to be published as soon as it's available (at or beyond the expected time)

- Including late data
- Where no data available from source system, don't publish
- Publication timeliness requirement to be monitored and the Business (including App Support) notified of breach

14. Data published to be available online for 2 years.

Non-Functional

15. Flexibility to choose/configure what data to acquire from iGMS Network Manager and specified Business Apps. (In iGMS Data Publication Hub scope).

16. Flexibility to choose/configure what data to publish

17. Flexibility to choose/configure frequency of publication

18. Flexibility to choose/configure recipients

19. The iGMS performance must not be degraded by the solution for this requirement. (In iGMS Data Publication Hub scope).

20. Future changes in data publication should not incur major performance and regression testing overheads

21. There must be no degradation of performance in the publication of existing IE reports. Transco have financial liabilities associated with the publication of some existing IE reports.

22. If they approve the UNC mod in April, Ofgem require solution implementation on 1st October 2006.

23. Environments will be required to support all project stages including Dev, OAT, UAT, Prod, Post Implementation testing (cut down version to allow for external Users

development testing). This will require an optimum number of discrete environments and subsequent management to ensure timely availability throughout project life-cycle.

24. It is anticipated that some Users will employ automated polling mechanisms that may affect access of data for other Users. The solution must actively negate any adverse affect on non-automated Users.

Service Level Agreement Requirements

25. The minimum SLA requirements are those of IE3 post- resilience project work.

26. In addition the following has been specified by the Business:

- Acceptable data loss = none
- Support for service = 24/7
- Time to recover = a maximum of 2 publication slots which equates to 30 minutes.

Appendix 10.b.i

Date	View count	Download count	Hit count
01-Sep-05	27	4643	4670
02-Sep-05	7	4354	4361
04-Sep-05	4	4394	4398
05-Sep-05	26	4599	4625
06-Sep-05	7	2465	2472
07-Sep-05	16	5597	5613
08-Sep-05	10	5694	5704
09-Sep-05	9	5479	5488
10-Sep-05	0	4972	4972
11-Sep-05	2	6176	6178
12-Sep-05	10	5350	5360
13-Sep-05	8	5398	5406
14-Sep-05	4	2594	2598
15-Sep-05	36	7037	7073
16-Sep-05	17	5296	5313
17-Sep-05	4	4768	4772
18-Sep-05	11	5563	5574
19-Sep-05	20	5123	5143
20-Sep-05	20	5728	5748
21-Sep-05	15	5372	5387
22-Sep-05	19	5575	5594
23-Sep-05	24	5180	5204
24-Sep-05	11	9966	9977
25-Sep-05	5	6238	6243
26-Sep-05	25	5514	5539
27-Sep-05	73	7322	7395
28-Sep-05	47	5637	5684
29-Sep-05	12	4403	4415
12-Oct-05	0	1059	1059
13-Oct-05	0	105	105
14-Oct-05	5	1218	1223
15-Oct-05	27	634	661
16-Oct-05	28	530	558
17-Oct-05	37	2600	2637
18-Oct-05	19	2762	2781
19-Oct-05	22	2831	2853
20-Oct-05	14	2899	2913
21-Oct-05	10	2828	2838
22-Oct-05	25	2957	2982
23-Oct-05	18	5079	5097
24-Oct-05	38	3152	3190
25-Oct-05	38	2821	2859
26-Oct-05	12	2560	2572
27-Oct-05	14	6508	6522
28-Oct-05	6	7557	7563
30-Oct-05	25	23517	23542
31-Oct-05	51	9673	9724
01-Nov-05	22	9495	9517
02-Nov-05	29	7696	7725
03-Nov-05	31	7131	7162

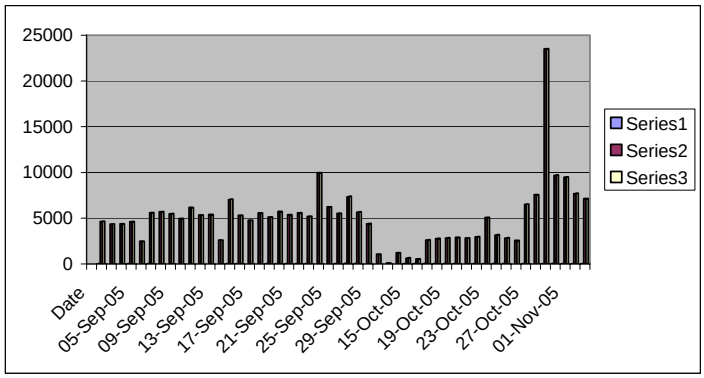

Appendix 10.c

Energy Watch Logical Architecture

1. Introduction

As a result of the DTI Information Initiative NGT has been tasked with publishing information regarding the gas transmission system via a public website. Most of this data will be sourced from iGMS, although other sources might be involved in the future. Several "categories" of data are already published via the Information Exchange (IE3) website, but there is now a requirement to extend this to include "near real time" publication of sub-terminal, onshore field and storage withdrawal flows.

The existing IE3 architecture may not be appropriate to meet the new requirements. This document explores the key requirements which impact on a solution architecture, sets out an appropriate logical architecture (with a number of options), and explores the possible relationship to IE3.

2. Key (Architecturally Significant) Requirements

A number of requirements critically dictate the system behaviour and the architecture(s) which can therefore be used. These are outlined below with current known values or working assumptions. Each of these must be confirmed, and the proposed architecture reviewed accordingly.

2.1 Data Rates and Volumes

The system will publish data from iGMS at five minute intervals. For each interval the data will be a snapshot of flow rates at each of around 40 metered points. The exact data content is TBC, but is estimates to be 20-100 bytes per reading, a total of 7-48 Kbytes/hour.

The system will retain up to two year's data online, a total of around 8.4 million readings, or less than 1GByte data volume.

2.2 User Numbers, Load and Performance

There are around 100 "expected" users from Shippers and other regular market participants. Public, casual and incidental use is unlikely to more than double this number.

Response times are TBC, but as a base assumption common queries should return within 15s, which is commonly regarded as a standard figure for internet applications.

The system is likely to be "polled" by users, rather than providing an event-driven notification system. Some of this polling is likely to be automated. Ideally each user

would poll once in every update cycle, generating a low transaction rate of around 0.7 queries per second.

However, in practice automated queries will run more quickly than this, particularly if certain users believe that they may gain competitive advantage from quickly accessing the latest information. It would therefore be reasonable to design the system to serve one query to each user within the standard response time. This is a rate of 13 queries per second.

It is, of course, possible for users to issue queries in excess of this level. However, NGT must learn the lesson from previous problems (e.g. Gemini) and manage the load rather than accepting any arbitrary and artificially high demand. A possible approach to this is described under "
3.2 Performance Management" below.

It is sensible to allow some headroom for application overheads, additional data or query types, and update transactions. Therefore *the system should be designed for 20 transactions (including queries) per second.*

There is no plan to undertake significant post-processing on the data after extraction from the source systems.

The maximum permitted latency from source system update to publication is five minutes.

Use of this system should not impact the performance of the iGMS or the existing IE reports.

2.3 Availability and Reliability

The system should be "continuously available", with no protracted planned outages. It is not an operationally critical system, but regulatory requirements mean that some financial penalties may be payable for any significant unplanned outage (TBC). However, given the five minute data change cycle an interruption of a 1-2 times this duration should be readily tolerable (TBC).

This is predominantly a query-only system with few or no user updates to data. Therefore there is no need to persist user session information through any failure, and simple "batch" transaction techniques should be adequate to protect update processes.

2.4 Security

The system will be a "Public access" system, with all data available to all users. There will not therefore be any requirement to authenticate users or manage identity information beyond that required for simple usage statistics.

2.5 Flexibility

If possible without significant extra cost the system architecture should allow for several expected dimensions of future flexibility:

- Use for publishing other data, possibly from different source systems and delivered via different interfacing mechanisms,
- Improvements to data latency and publication rates. The current assumption is that 1 minute would be a reasonable minimum target for each.
- Improved “speed to market”, i.e. the ability to quickly publish similar data without the need for a protracted system development. In particular, it must be possible to change data sources and publication formats without a major regression test exercise.

The data format should allow the users to easily review data on-screen and also to download and further process the same data.

2.6 Delivery timescale

The latest date for delivery is July 06, but Ofgem would prefer something sooner, e.g. Oct 1st 2005.

3. Logical Architecture

3.1 Overview

The following diagram summarises the proposed solution, which is a standard 3-tier web application using modern component technologies:

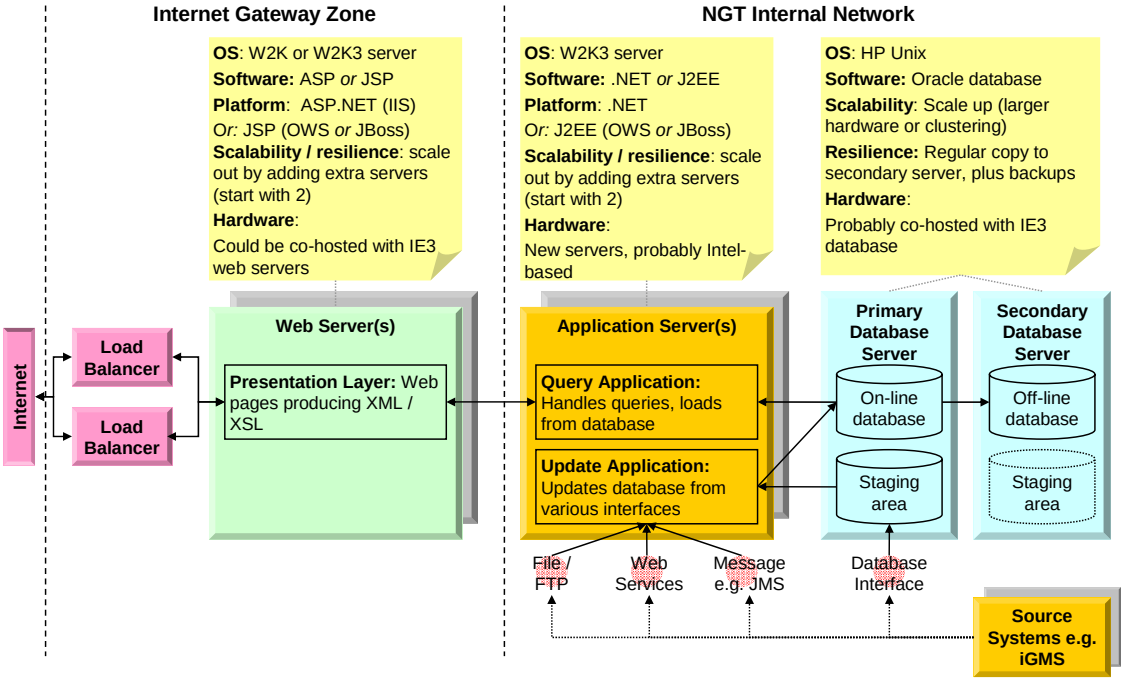

The web server focuses on presenting the web pages to the user, navigation logic, and output formatting. It connects to the application layer in the middle tier, but does not connect directly to the database. The “presentation” application will use a combination of static HTML pages, Active Server Pages (or Java Server Pages), and XML. Query results will be formatted as XML, which can then be either presented to the user using XSL stylesheets, or downloaded and further processed if required.

The middle tier will be written using either Enterprise Java (J2EE) or Microsoft .NET component technologies (to match the presentation layer). There will be two “applications”:

- The query application. This will receive query requests from the front-end and return responses. Where necessary this will interact with the database to extract appropriate data, but in many cases the response will be constructed from data held in active objects.

- The update application. This will service the incoming interface(s), processing incoming data and submitting updates to the database. A variety of interface mechanisms can be supported within a consistent processing architecture. Oracle-based interfaces may deliver data directly to a “staging” area – the update application will treat this area in the same way as any other incoming interface, processing the data and managing updates to the main on-line database.

Both applications will share a “data access layer” which will manage interactions with the database. That will be a relatively “dumb” Oracle database hosted on a primary server, with a warm-standby secondary server maintained from the primary by a background data mirroring process.

The use of an Oracle database is a sensible “default” option, but not the only one. Data volumes are sufficiently low that other alternatives, including file-based storage, could also be considered.

3.2 Performance Management

The expected transaction rate suggests a multi-tier system, with an intermediate application layer managing active data and database connections independent of the front end. While it is possible that the performance requirement might be met from a single server at each layer, the availability requirements dictate at least two. Therefore the front and middle tiers are designed to use two or more relatively small servers, and throughput can be increased by “scaling out”, adding more identical servers at the appropriate layer.

The database server should be lightly loaded in this design, and the databases may sensibly be co-hosted with other Oracle services on appropriate large shared servers. In the unlikely event it is necessary increased capacity can be achieved by scaling up or clustering those servers.

3.2.1 Demand Throttling

NGT must learn from previous experience and manage user demand so that the system is not “thrashed” by unlimited automated polling of the web pages. The requirement to publish the data via simple public web pages prevents a true event-driven approach (e.g. publishing data to message queues). Failing to properly manage demand will destroy any “level playing field”. Therefore the application must impose a degree of demand management itself. Various techniques are possible. The following summarises one simple possibility:

- Users access the data via a “gateway” page. This includes the time of last publication, and a link to the full query page. The gateway page is a “static” HTML page (updated once every five minutes) so that it can be polled frequently with impunity.

- The users can only get the current data by waiting for the gateway page to change, then following the link. Re-polling a previous query page will only deliver old data.

- The link contains an “encrypted” element which changes at each publication, so that users cannot “predict” the address of the new data and repeatedly poll it.

3.3 Availability and Reliability

The proposed architecture should require very few planned outages, and any outage should be short, limited only to the time taken to reconfigure the system’s connectivity (i.e. which servers are active at each layer) and restart appropriate processes.

The database can be backed up from the off-line copy on the secondary server. There is no persistent data on any other server, so backups need only be taken occasionally, probably with the system in use.

There is no single point of failure. Any component can be taken out of service for maintenance or updates. Most updates can be applied to one node first, (while the others continue to support the live service) then the other(s). The proposed component technologies both also support “hot” updates to the application components if that is appropriate.

3.4 Data Management

The database in this architecture is used to “persist and populate” object data, rather than to serve every query. It also has a low projected total volume. This allows it to have a very simple structure, and to be managed with simple tools.

More complex management software such as Oracle Warehouse Builder should not be required, but could be introduced later if it becomes necessary to manage more complex database interfaces.

There is probably no case for a complex, “heavyweight” data access mechanism serving the application layer. To reduce complexity the detailed design should consider adopting a lightweight persistence framework (such as Hibernate for a Java-based design).

3.4.1 Non-database option

The projected data volumes and data update activity are relatively low. The database should be co-located with others on suitable shared servers. If there are any barriers to this approach it may be appropriate to consider alternatives such as an XML file-based store. Adoption of an abstract persistence framework such as Hibernate would allow the data store to be changed independent of the rest of the application, which might also be useful for development and test purposes.

3.5 Interfacing

The proposed architecture would support a variety of interface mechanisms, via interface-specific handlers “plugged in” to a common set of components preparing and handling updates to the database.

The best mechanisms would be file-based (including FTP), message-based (e.g. JMS) or web-service (SOAP) based, allowing the source system to be loosely coupled with the Energy Watch system, and allowing updates to queue up for later processing in the event of a system or network outage or slowdown. JMS is probably the best option given the Unix/Java basis of key source systems.

The architecture also supports interfaces which deliver data directly to a “staging area” on the Oracle database. This may appear easier where the source system also has an Oracle database, but has a number of implications:

- The source system must either produce a stream of updates (e.g. via database update triggers) or allow regular queries. Either may have some impact on the source system.

- Direct database updates require both systems to be available concurrently, which may not be the case where the source system has a very high availability requirement (e.g. iGMS). The option is to use Oracle-based queuing (AQs), but this increases complexity.

- Discussions about how best to serve the data from iGMS are ongoing and a final recommendation will follow.

3.6 Platform Choice

This architecture could be built readily using either J2EE or Microsoft technologies. The platform choice will depend largely on supplier skills, and the proposed implementation approach.

The only significant technical indicator for J2EE is if iGMS uses it and can easily drive JMS, although it would be possible to bridge an MS-based system to JMS if required. The other interface options (database, file and web services) are all platform-independent. Performance might also be a discriminator (small MS-based systems tend to out-perform J2EE) but this is unlikely to be critical. A final decision is not needed until issues of approach and sourcing have been resolved.

4. Relationship to IE3

IE3 is designed to deliver complex Business Objects reports on a relatively long cycle (currently every 10 minutes). The core application servers are quite heavily loaded, and there are already performance issues. There are also issues The application architecture is not suited to delivering smaller data updates on a short cycle - the possibility of changing the IE3 update cycle to 5 minutes or shorter has been reviewed by a third party consultant which and did not think this could be easily achieved. around the software versions currently installed, some of which are now out of support.

The IE3 web servers and load balancers are relatively lightly loaded, and the database is to be rehosted on new HP Unix servers which should also then have some spare capacity.

The Energy Watch application cannot be delivered “under IE3”, as the application architectures are quite different. However, there are some opportunities for infrastructure rationalisation. The physical architecture is strictly not an NGT issue but it is possible to propose an obvious solution which will minimise any new infrastructure:

- The Energy Watch presentation layer will be co-hosted on the same web servers as IE3, sharing the same load balancing and (if required) user management services,

- The database will be co-hosted on the same database servers,
- The application layer will be hosted on new servers dedicated to Energy Watch. For the production system these will be two small-medium servers, probably Intel-based.
- A similar approach can be taken for test environments.

5. Implementation Strategies

The proposed Energy Watch application should be relatively simple to build, and NGT should carefully consider how to structure and resource the project to minimise initial and lifetime costs. While the iGMS data extract mechanism will have to be built through a formal process to manage its impact, the rest is ideally suited to an agile, iterative and incremental development. It is this paper's recommendation that such an approach be actively considered.