Fixed Telecoms Switching

Technical Study

20 January 2021

DOCUMENT VERSION 1.0

Version History

Version Number	Implemented By	Revision Date	Approved By	Approval Date	Description of Change
1.0	Dave Fitton	20/1/2021	Michael Dargue	20/1/2021	First issue

Copyright

Cartesian Ltd. Registered in England and Wales. Registered Number: 03230513

Registered Office Address: Descartes House, 8 Gate Street, London WC2A 3HP United Kingdom

Table of Contents

List of Figures

List of Tables

1. Executive Summary

Ofcom is considering options for a new switching process to enable residential customers to switch their fixed voice and broadband services simply and reliably.

This technical study reviews the business process and systems architecture of two industry proposals for how a new process could work, namely Option X and Option Y. The scope of the study covers technical risks in the proposed business processes and supporting systems.

The defining feature of Option X is a switching code that the customer obtains from the Losing Communication Provider (LCP) and provides to the Gaining Communication Provider (GCP). The switching code provides a mechanism to identify the services to be switched.

Option X proposes a hierarchical hub and spoke architecture, with the Access Providers interfacing to a central hub and providing switching services for their respective wholesale customers. By default, communications between the retail Communication Providers (CPs) and the Hub flow up and down the CPs' supply chains; however, Option X suggests that large CPs could also connect directly to the Hub if there was value in this approach.

In Option Y, the customer requests a switch by providing the GCP with their name, address, and the identity of their LCP. These details may be supplemented by additional information if required. The GCP then communicates with the LCP via a central hub to identify the services to be switched.

In Option Y, the retail CPs interface directly with the Hub and are not reliant on their supply chain to initiate a switch; the supply chains are notified to enact the switch afterwards.

A further difference between the two options is in the communication of Switching Information (SI). In Option X, this is provided by the LCP to the customer when the customer requests the switching code.

In Option Y, the information is communicated indirectly from the LCP to the customer – either via the Hub (Option YHUB) or via the GCP (Option YGP). Option Y proposes that the Hub and GCP would be able to provide SI in durable form by email, text message (with a link to a PDF), or letter for call centre, and hard copy for in-person transactions. For online channels, the SI would be displayed as inline content on the GCP's website.

The study identifies specific risks for both options throughout the switching journey.

In both options, the systems operated by the GCP, LCP or Hub may be unavailable for planned or unplanned reasons. In most cases, an unavailable system would cause a temporary delay to a customer's switching journey and this may require the customer to resume or reattempt the switch once the issue is resolved.

It is not possible to quantify the risk of an individual system being unavailable or determine how quickly a system could be recovered from failure. Much of this will depend on design and implementation choices. Both options are well thought out technically and could be implemented successfully.

Both options include a requirement for a Hub to act as a central point to support customer switches. The availability of the Hub is therefore a key design consideration for both options.

For Option X, the dependency of retail CPs on their supply chain for communications with the Hub increases the risk that a customer may be forced to interrupt their switching journey. This risk increases with the length of the supply chain but may be mitigated by wholesalers and Access Providers (APs) adopting higher availability systems. Option X also provides for large retail CPs to connect directly to the Hub.

Specific to Option Y, and aside from the risk of system failure, is the risk that the GCP will encounter issues when seeking to identify the customer's existing services with the LCP. This is most likely to occur due to a problem in address matching. The Option Y team estimates a success rate of 99% for address matches using a Unique Property Reference Number (UPRN) and an overall success rate of identifying customer's services above 96%.[1] The 99% estimate seems reasonable based on Cartesian's previous experience in working with large UK CPs; however, we have not sought to validate this with CPs within the scope of this project.

In Option Y, where an initial match cannot be made, the customer would need to supply additional information such as their landline phone number or account number. If this is unsuccessful a fallback mechanism will be required, however this is expected to be a very rare situation.

The study also compares the two options in terms of their relative complexity, security, and ability to cope with the expected volume of transactions.

In terms of systems architecture, the two options have fundamental differences which each introduce their own complexities. In Option X, the GCP and LCP independently communicate with the Hub. In Option Y, the GCP and LCP communicate with each other via the Hub in near real time.

The independence of the GCP and LCP interactions in Option X should make it easier to isolate faults between the gaining and losing parties. However, this benefit needs to be balanced against the use of the supply chain in Option X for communication with the Hub. Involving the supply chain adds additional parties and interfaces which adds complexity to Option X.

Option Y is found to be more complex in the indirect communication of the SI.

The technologies proposed for both options have been available for many years and are well understood. As such, the technology does not introduce any specific or unusual risks.

From a security perspective, neither solution is materially better or worse.

The transaction volumes are well within the capability of modern IT systems and do not introduce any further risks for either option.

Finally, the study considers potential implementation issues. Both of the options could be introduced with a period of parallel running to enable switching and porting requests that are already in progress to complete under the existing processes. Both options could also be phased in over time, however the benefits of this may be outweighed by the potential customer confusion.

In summary, neither of the two options has a materially higher level of risk on paper. They each have specific failure modes which will need to be considered in the design and implementation phases.

[1] OTA, Y Group answers to Ofcom questions: "*What is the expected success rate of matching against the mandatory info (name, address, postcode, LP name, services)? Above 96%*". Also, footnote: "*We estimate each of the 4 factors (UPRN, customer name, retailer name and services) as having at least 99% matching for consumer customers. Overall this give a 96% success rate*"

2. Introduction

Ofcom is considering options for a new switching process to enable residential customers to switch their fixed voice and broadband services simply and reliably.

An industry working group convened by the Office of the Telecommunications Adjudicator (OTA) has developed two proposals for how a new process could work.

This technical study reviews the business process and systems architecture of the proposals, namely Option X and Option Y. Further information on each proposal is available in the reference documents supplied by Ofcom and created by industry experts. See Reference Documents section on page 25.

3. Overview of Proposed Options

3.1 Option X Overview

The architecture proposed by Option X consists of a central Hub which is used by the retail CPs to coordinate switching requests. Option X assumes that retail CPs do not need to connect directly to the Hub. As shown in Figure 1, the connection can be made via the APs and other parties in the CP supply chains.[2]

In Option X, the Hub performs the following primary functions:

- Generates and issues switching codes against LCP services
- Tracks the status of switching codes and allows CPs to check code status
- Enables GCPs to use switching codes to initiate a switch
- Notifies the LCP supply chain of switch requests
- Provides a mechanism for the GCP to cancel a switch before the Point of No Return (PONR)

Figure 1. Architecture of Option X with Hub at Network/VoIP Provider Level

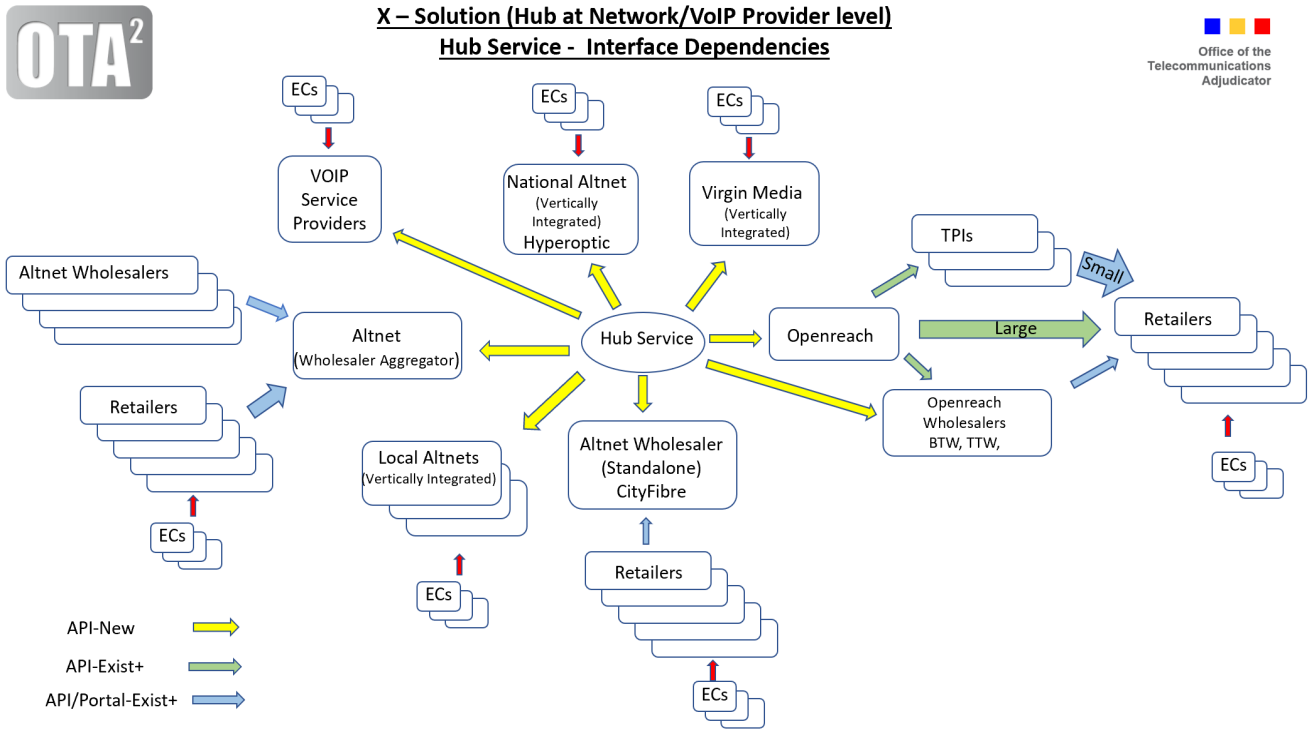

Source: OTA

[2] Note that Option X does not preclude additional entities connecting to the Hub if there is value in doing so.

3.2 Option Y Overview

The architecture proposed by Option Y also has a central Hub which is used by the retail CPs to coordinate switching requests. In Option Y, all retail CPs connect directly to the Hub, with the exception of CPs that are served by Third Party Integrators (TPIs). This is illustrated in Figure 2.

In Option Y, the Hub performs the following primary functions:

- Provides the LCP and GCP with a messaging service for them to identify the services to be switched, initiate a switch request, and cancel switches before the PONR
- Generates and issues switching order references
- Provides an audit trail for communications between the LCP and GCP

There are two variants of Option Y. In Option YHUB, the Hub is also responsible for communicating SI from the LCP to the customer. In Option YGP, this function is performed by the GCP.

Figure 2. Architecture of Option Y

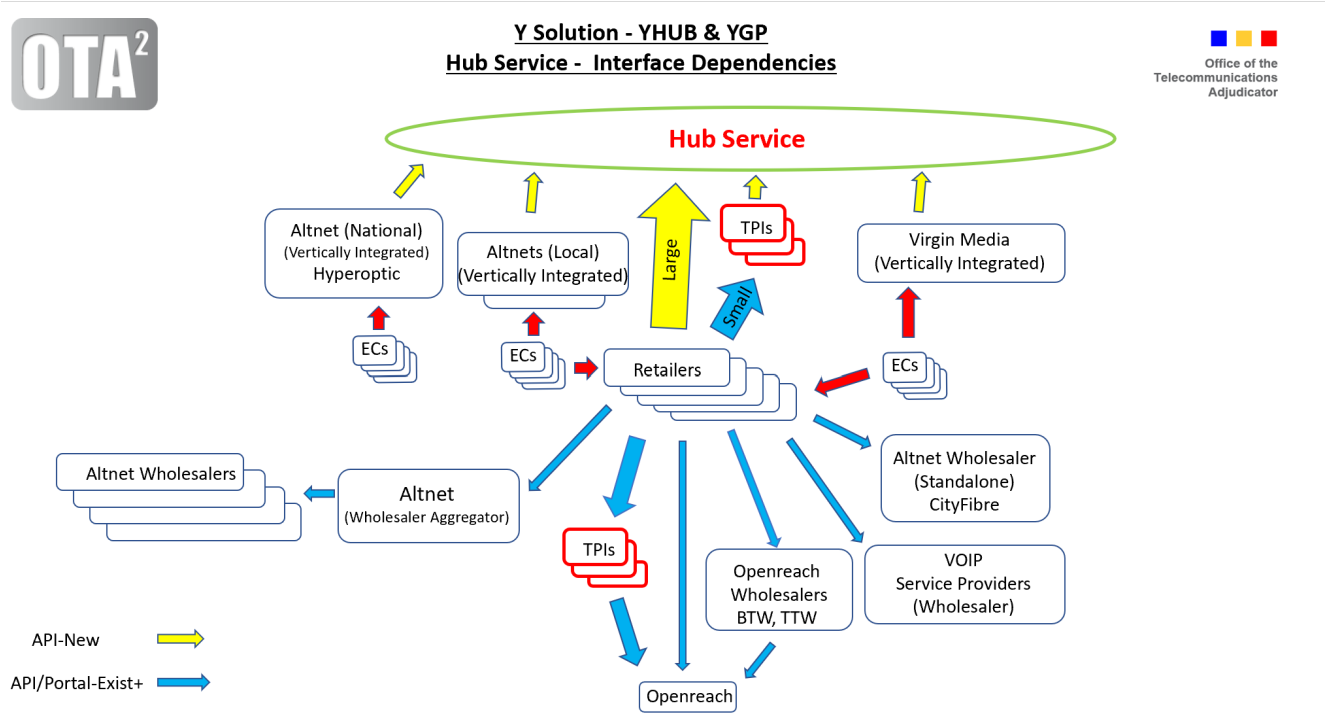

Source: OTA

4. Business Processes and Risks

In the following sections the essential business functions of the switch process are considered:

- Identifying the services to be switched
- Informing the customer of switching implications
- Initiating the switch
- Handling in-flight changes

For each function, the approach proposed by each option is summarised, followed by an assessment of perceived risks. Where appropriate, observations on specific channels are made.

For simplicity, references to retail CP systems are taken to include any TPI systems that the retail CP is dependent upon.

4.1 Identify Services to be Switched

Correctly identifying the services to be switched is important to avoid erroneously switching either (i) services belonging to a different customer, or (ii) the incorrect set of services for the right customer.

Note: Residential fixed line voice suppliers have obligations around making caller location information available to emergency services, which adds extra importance to address accuracy.

4.1.1 Proposed Approach – Option X

The customer contacts the LCP and goes through the standard Identification and Verification (ID&V) process. Option X proposes that call centre and online channels are supported as a minimum for switching requests, but LCPs may choose to offer additional options. The ID&V process will be specific to the channel.

To initiate the switch, the customer selects which services they would like to switch (broadband, fixed voice, or both). For fixed voice, the customer also indicates if they would like to port their phone number.

The LCP identifies the relevant services and communicates the SI to the customer (see Section 4.2.1).

The LCP then requests a switching code which is generated by the Hub. By default, Option X proposes that the LCP communications are routed to/from the Hub via the LCP's supply chain. Optionally, the LCP may interface directly with the Hub.

Having obtained the switching code, the LCP then provides it to the customer. The switching code identifies the service set that the customer indicated they would like to switch.

The switching code should be returned to the customer within a short period of time. Option X proposes the code should be returned to the customer within 60 seconds, with a design target of 10 seconds[3].

[3] OTA

Later in the process, the switching code is used when the customer places an order with the GCP. If the code does not match the services that the customer asks to switch, the request will be rejected.

4.1.2 Proposed Approach – Option Y

The customer contacts the GCP. The GCP collects details from the customer, including the service address, name of account holder, contact details for the account holder, the LCP identity, and the services to be switched.

The GCP uses the information gathered from the customer to identify the service to be switched with the LCP. The GCP and the LCP communicate via the Hub.

To identify and validate the services to be switched, the LCP will require – as a minimum – the customer service address, the customer name, the LCP identity, and the services to be switched. UPRNs will be used to identify addresses between CPs.

If there is a full match, then the switching request is accepted. If a full match cannot be made, the LCP informs the GCP via the Hub. In this case, the GCP can reattempt the match with additional information such as the LCP account number or phone number.

Once a full match is made, the Hub creates a switch order reference that identifies the services to be switched. The switch order reference is forwarded to the GCP.

If two services are to be consolidated onto the GCP from different CPs, then the GCP will need to validate customer address and service details with both LCPs. This is in order to obtain two switch order references, one for each of voice and broadband from the respective LCPs.

4.1.3 Risks Identified

Risk 1. ***Selected channel not available (Option X and Option Y)***

There is a risk that the channel selected by the customer is not available due to planned or unplanned technical reasons (i.e. not just unavailable due to being outside of the channel trading hours). For example, a web portal may be offline or the support systems in a call centre may be temporarily unavailable due to a technical fault.

Depending on the nature of the problem, the customer may be able to proceed with their switch through an alternative channel. Otherwise, the customer would need to reattempt the switch at a later time.

In general, CP channels have good availability – sufficient to operate business as usual activity in customer acquisition and support. If a temporary channel failure occurred, the impact would be limited in scope, affecting only those customers seeking to leave a specific CP via a specific channel at the time of the failure. The time to recover from such a failure will depend on the specific nature of the issue; however, in most cases this should take no longer than a few hours.

Risk 2. ***Back-End System Failure (Option X)***

There is a risk that one or more of the back-end systems are unavailable due to planned or unplanned technical reasons. This would prevent the LCP from obtaining a switching code. The customer would need to try again at a later time when the systems are back online.

To obtain a code, the LCP systems, the Hub, and any intermediary systems in the supply chain must all be available.

If an intermediate system is not available, that will affect the customers of all downstream LCPs that were reliant on that system. This could affect a large segment of the market if, for example, a system of a large AP failed.

It is assumed that the Hub would be designed for high availability, given the high dependency on this central system. The broad impact of failures in intermediary systems would also need to be considered in the design phase. Greater investment in these systems (e.g. providing full redundancy) may be required to achieve industry availability targets. This is explored further in Section 5.2.2.

Risk 3. *Back End System Failure (Option Y)*

There is a risk that one or more of the back-end systems are unavailable due to planned or unplanned technical reasons. This would prevent the GCP from identifying the services to be switched with the LCP. The customer would need to try again at a later time when the systems are back online.

To identify the services to be switched, the LCP systems, the Hub, and the GCP systems must all be available.[4] It is assumed that the Hub would be designed for high availability, given the high dependency on this central system.

Risk 4. *Customer Details Fail to Match (Option Y)*

There is a risk that the information provided by the customer fails to match the records held by the LCP. Should this occur, the customer would be unable to proceed with a switch until a positive match is obtained. Instances where addresses fail to match will only impact a single switch attempt.

In Option Y, the GCP will seek to match the customer against the LCP records using a combination of customer address, surname, and LCP name.

Option Y proposes that addresses are exchanged as UPRNs as this will help to identify individual premises within a larger property. UPRNs provide a good reference for addresses, however a low level of address matching issues is likely to be unavoidable.[5]

The Option Y team estimate a better than 99% matching rate to resolve the correct UPRN from a given address and postcode.[6] The estimate seems reasonable based on Cartesian's previous experience in working with large UK CPs; however, we have not sought to validate this with CPs within the scope of this project.[7]

The Option Y team also estimate a matching rate of better than 99% for customer name and LCP identity, given some basic data normalisation. Again, this estimate seems reasonable based on

[4] For retailers that are dependent on TPIs, this includes the TPI systems.
[5] Problems relating to obtaining UPRN and address data occur most frequently for new properties. Whilst new-build properties are unlikely to be encountered in a switching scenario, situations where an existing property is split into flats may cause problems. For example, the LCP may hold the address of the parent property and this would fail to match the UPRN of a specific flat. Allowing for parent-child property relationships in the UPRN matching should catch many of these issues. (A useful discussion of these issues can be found in Ofcom's "Connected Nations 2019 Annex A: Methodology".)
[6] See footnote 1.
[7] Although historically poor, address data quality has been improved by CPs aligning to third party databases for address management (i.e. Postcode Address File (PAF) for mailsort discounts, etc.) and aligning to addresses of Access Providers (i.e. Openreach). For cable networks, the historical consolidation of cable operators under Virgin Media drove improvement in address data quality as systems were standardised.

Cartesian's experience, but we have not sought to validate this with CPs within the scope of this project.

Taken together, the Option Y team estimate a success rate better than 96% for matching on name, address, postcode, LCP name, and services.[8] For the remaining 4%, Option Y suggests using additional detail (e.g. account number or phone number) to resolve matching issues. The LCP will need to be able to incorporate these additional details into the matching process.

Should the matching issue persist, then a fallback mechanism will be required, however this is expected to be a very rare situation.

Risk 5. *Customer Details Incorrectly Matched (Option Y)*

If a customer's address is incorrectly matched, there is a risk that this could trigger the switch of the wrong service (i.e. an erroneous transfer).

Option Y protects against this risk by requiring that the customer's surname and LCP identity must also match. The likelihood of getting a false positive when using a combination of address, surname, and LCP is extremely low.

Risk 6. *Incorrect Service Identified (Option Y)*

If a customer has more than one service (e.g. two phone lines) with the LCP at a given address, there is a risk that the incorrect service is matched.

A full match will require identifying the service within the address by use of the optional details suggested in the Y process, e.g. the phone number or other identifiers such as the serial number of the LCP's customer premises equipment (e.g. Fibre-To-The-Home Optical Network Terminal). The customer may need to get this information from the LCP or from reference labels in their home where the service is presented.

It is worth noting that having more than one residential service at an address is uncommon. Where this is the case, the customer may be more familiar with technology than an average residential customer.

4.2 Inform Customer of Switching Implications

Switching services gives rise to SI. This information must be communicated to the customer in a durable form and in a timely manner. SI also needs to be made available to the customer during a live interaction (online/phone).

The LCP compiles the SI for the customer based on the services being switched, the customer's contract status, and any associated services.

Some CPs may have the capability to generate SI on demand; alternatively, an up-to-date set of SIs could be maintained (e.g. calculated overnight) and stored for retrieval when needed.

[8] See footnote 1.

4.2.1 Proposed Approach – Option X

In Option X, the SI is compiled by the LCP in response to the customer's request to switch services. The LCP will present or explain the implications to the customer during the live interaction and will provide the SI in a durable form.

Options for communicating SI in a live session will be determined by the channel the customer has selected. If the customer is talking to an advisor in the call centre, then the SI can be explained verbally. For online transactions, the SI can be presented in the LCP customer self-care portal or a mobile app.

Options for also providing in a durable medium will depend on the channel and the customer's preference, and would include email, download from a website, or letter (if requested).

At a later point, when the customer initiates the switch with the GCP, the LCP will need to confirm the SI in a durable form with the final bill. The LCP will need to recompile the SI at this point in case there have been any changes since the LCP interacted with the customer. Communication of SI with the final bill will follow existing, business-as-usual processes.

4.2.2 Proposed Approach – Option Y, YHUB variant

In the YHUB variant of Option Y, the SI is sent to the customer by the Hub in response to a request from the GCP. The LCP compiles the SI and sends the information back to the Hub. The Hub then forwards the SI to the customer.

For call centre and in-store channels, the Hub will communicate the SI in durable form using customer contact details provided by both the GCP and LCP. Option YHUB proposes that SI may be sent by email, text message, or letter. A text message can include a link to a reference document, for example in PDF format, which provides more information than is accessible via a text.

For online channels, Option YHUB proposes that SI are displayed as inline content, using a mechanism similar to the 3D Secure system adopted for online card payments.[9] The SI content would be served by a web server at the Hub and would appear within the GCP's webpage; however, the GCP would not be able to access the LCP content. The information would be retrievable later by the customer.

In addition to communicating SI via the Hub, the LCP will also send the SI directly to the customer with the final bill following existing, business-as-usual, processes.

4.2.3 Proposed Approach – Option Y, YGP variant

In the YGP variant, the Hub requests the SI information from the LCP as for YHUB above. However, rather than forwarding this information directly to the customer, the Hub forwards it to the GCP. The GCP can then communicate the SI to the customer in a live session and provide the information in a durable form.

In addition to communicating SI via the GCP, the LCP will also send the SI directly to the customer with the final bill following existing, business-as-usual processes.

[9] HTML inline frame. For more details, see https://www.w3.org/TR/html4/present/frames.html#h-16.5

4.2.4 Risks Identified

Risk 7. ***Delay in generating SI information (All Options)***

Under normal conditions, timely access to the SI details should not pose a problem as the risk of system failure for an individual CP is very low.

In the event that an LCP experiences a systems failure, customers seeking SI from the affected CP would not be able to obtain their SI. The customers could resume/restart their journey once the issue is resolved.

Risk 8. ***Failure to Communicate SI to Customer (Option YHUB)***

Communicating the SI to the customer requires messages to be passed between CPs and the Hub, and for these to be processed in near real time. For YHUB the following steps must all complete successfully:

1. GCP sends SI request to the Hub
2. Hub contacts LCP requesting SI
3. LCP sends SI information to the Hub
4. Hub sends the SI information to the customer through selected channels

There is a risk that the GCP, Hub or LCP systems may not be available due to planned or unplanned technical reasons. In each case, this would delay the SI information reaching the customer.

The scope of impact will depend on which system(s) are unavailable:

- If the GCP systems are unavailable, only customers trying to switch to that CP will be affected.
- If the LCP systems are unavailable, only customers trying to switch from that CP will be affected.
- If a TPI's systems are unavailable, the customers of CPs reliant on that TPI will be affected.
- If the Hub is unavailable, all switches will be affected.

Whilst the possibility of any given interface being unavailable is low, the requirement to have LCP, GCP and Hub interfaces continuously available multiplies the probability of a failure. This is explored further in Section 5.2.2.

If the SI is not available, then the customer will be unable to proceed with their switch order. The customer could resume/restart their journey once the issue is resolved. It is unlikely that the system(s) will be unavailable for significant periods.

Risk 9. ***Failure to Communicate SI to Customer (Option YGP)***

Similar to YHUB, for the SI information to be generated, a sequence of messages must be processed in near real time. For YGP the following steps must all complete successfully:

1. GCP contacts the Hub requesting SI
2. Hub forwards request from GCP to LCP
3. LCP sends SI information to the Hub
4. Hub forwards the SI information to the GCP
5. The GCP presents/sends information to the customer.

Again, there is a risk that the GCP, Hub or LCP systems may not be available due to planned or unplanned technical reasons. In each case, this would delay the SI information reaching the customer.

The scope of impact will depend on which system(s) are unavailable, in the same way as for YHUB above.

In the event of a failure, the customer could resume/restart their journey once the issue was resolved. It is unlikely that the system(s) will be unavailable for significant periods.

4.3 Initiating the Switch

To initiate the switch, the GCP needs to inform the LCP which services are being switched, and when the switch will occur. Both the GCP and LCP then need to work with their respective supply chains to enact the switching activities, i.e. provision the new service and cease the old service.

The switching activities will vary depending on whether the switch is inter- or intra-network. If porting is requested, then this will also need to be addressed.

4.3.1 Proposed Approach – Option X

Before contacting the GCP, the customer must have already obtained a switching code from the LCP. The switching code provided by the LCP will relate to the specific service(s) being switched:

- Fixed Voice only;
- Broadband only; or,
- Fixed Voice and Broadband.

To initiate a switch, the customer provides their customer information (service address, etc.) and the switching code to the GCP. If two services are to be consolidated from different LCPs then two switching codes will be required, one for each service to be switched.

The GCP validates the switching code(s) against the Hub to check they have not expired, been claimed, or revoked. The GCP also confirms whether they can provide the service set associated with the switching code at the given address.

Once all the checks are complete, and the customer is ready to proceed, the GCP will agree a proposed switching date with the customer and complete the sale. The GCP will then engage with their supply chain to enact the switch.

The switching request is passed down the GCP supply chain to the Hub, using the switching code to identify the services to be switched. The Hub then notifies the LCP supply chain, and the notification is passed up through the supply chain to the LCP.

4.3.2 Proposed Approach – Option Y

The process for initiating a switch order is the same for both the YGP and YHUB variants.

The GCP sends the switch order to the Hub with the requested switch date. The GCP uses the switch order reference to identify the services to be switched.

Communications to coordinate the switch occur directly between the retail CPs via the Hub. The GCP and LCP then independently liaise with their supply chains to enact the switch activities.

4.3.3 Risks Identified

Risk 10. ***GCP unable to Validate Switching Code***

There is a risk that the systems required to validate the status of the switching code may not be available due to planned or unplanned technical reasons. This includes the systems of the GCP, Hub and any intermediaries between the GCP and the Hub. In each case this would delay the GCP from placing the switch order.

As explained above, the GCP and LCP may be reliant on intermediary firms in their supply chains for communicating with the Hub. These Hub integration services are expected to be timely and reliable, but the effect of daisy-chaining systems will compound the risk of downtime.

If the GCP systems are unavailable, that would affect all switch orders for that CP. If a wholesaler or AP's systems are unavailable, that would affect switch code validation from multiple GCPs. If the Hub is down, that will affect all validation attempts.

The GCP will be unable to proceed with the customer's request until the switching code is validated. The customer would need to try again at a later time when the systems are back online.

Risk 11. ***Communication Failure in Switch Order***

There is a risk that the systems required to communicate the switch order may not be available due to planned or unplanned technical reasons. For Option Y, this includes the systems of the GCP, LCP, and the Hub. For Option X, this additionally includes the systems of any intermediaries between the CPs and the Hub. In each case this would delay the GCP from placing the switch order.

If the GCP or LCP systems are unavailable, that would affect all switch orders for that CP. If the Hub is down, that will affect all switch orders. Specific to Option X, if the systems of a wholesaler or AP are unavailable, that would affect switch orders from multiple CPs.

It is unlikely that systems will be unavailable for significant periods. Order placement could resume once the failure has been resolved.

The customer would not necessarily need to be aware of this back-end issue. Note, however, that in Option Y, the switch order is placed by the GCP during a single interaction with the customer. A system failure relating to this risk may therefore also cause earlier stages of the customer journey to fail, as described in the preceding sections.

4.4 Handling In-flight Changes

Common to both options, changes can be made to the order up until the agreed PONR. Typically, these changes are to reschedule an order, or for the GCP to cancel the order for valid reasons.

There is also the potential for problems to be encountered beyond the PONR. In these circumstances, it may be necessary to restore the LCP service.

Neither of the options provides for an emergency restoration process. It is recommended that industry further explore the value of an emergency restoration during the design phase.

4.4.1 Proposed Approach - Option X

Option X provides for the GCP to cancel an in-flight order by sending a cancellation request to the Hub. The Hub records the change in state and notifies the LCP. All communications flow via the CP's respective supply chains.

Although there is no explicit option to delay a switch in the Option X proposal, it is understood that this could be achieved by cancellation of a request and resubmission with a new date. Alternatively, the design could be augmented to accommodate this use case.

4.4.2 Proposed Approach – Option Y

In Option Y, communications between the GCP and LCP are passed via the Hub. This communication mechanism would be used to handle cancellation or rescheduling of an in-flight request.

4.4.3 Risks Identified

Risk 12. ***Communication Failure for In-flight Change***

This risk is similar to the risk for a Communication Failure in Switch Order (Risk 11).

The consequence is that messages between the GCP and LCP could be temporarily delayed. Communications between the CPs would resume when the systems issues were resolved.

If there was an urgent need (e.g. to avoid a customer losing service), it is possible that CPs may seek a temporary workaround, i.e. communicating directly with each other and their respective supply chains.

4.5 Summary of Technical Risks

Risks identified in this section are summarised in the table below with the expected frequency, scope, and impact. Frequency is given in relative terms as, for most risks, this will depend upon design and implementation choices. (See Section 5.2 for further discussion on system availability.)

For Option Y, in addition to the risks relating to systems availability, there are also risks relating to the matching of customer details when identifying the LCP services. The most significant of these risks are (i) that a customer's details cannot be matched, and (ii) that the correct service cannot be identified where there is more than one service. An exception management process is likely to be needed to resolve specific customer cases.

Table 1. Summary of Risks

ID	Description	Option	Frequency	Scope	Impact
Identify Services to be Switched					
1	Selected channel not available	Both	Very Low	Customers using specific channel for that CP	The customer will need to try again later or use a different channel
2	Back end systems not available when attempting to obtain a switching code	X	Low	Ranging from single CP to all CPs depending on which system is affected	Depending on the length of outage, the customer may need to try again later
3	Back end systems not available when attempting to identify LCP services	Y			
4	Customer details fail to match LCP records	Y	Low[10]	Individual customer	The customer would need to validate the details and an address exception process may need to be invoked
5	Customer details incorrectly matched against LCP records	Y	Extremely low	Individual customer	Erroneous transfer if not detected, e.g. in SI comms
6	Incorrect LCP service matched (for same customer, same address)	Y	Very low	Individual customer	Erroneous transfer if not detected, e.g. in SI comms
Inform Customer of Switching Implications					
7	LCP systems issue causes delay in generating SI	Both	Very low	Customers for a specific LCP	Depending on the length of delay, the customer may need to try again later
8	Back end systems not available for communication of SI	YHUB	Low	Ranging from single CP to all CPs depending on which system is affected	Depending on the length of outage, the customer may need to try again later
9	Back end systems not available for communication of SI	YGP			
Initiating the Switch					
10	Back end systems not available for validation of switch code	X	Low	Ranging from single CP to all CPs depending on which system is affected	Depending on the length of outage, the customer may need to try again later

[10] Estimated by Option Y to be less than 4%. See footnote 1.

11	Back end systems not available for communication of switch request	Both	Low	Ranging from single CP to all CPs depending on which system is affected	The switch request would be delayed. Does not necessarily impact customer
Handling In-Flight Changes					
12	Back end systems not available for communication of switch cancellation	Both	Low	Ranging from single CP to all CPs depending on which system is affected	Switch cancellation would be delayed. Does not necessarily impact customer

5. Systems Architecture

This section reviews the technical systems architecture and related comments on each proposal.

Both proposals offer architecture options that show a good technical understanding and propose a similar core architecture consisting of a microservices-based hub with a messaging and notification configuration.

Each option presented is viable from a systems architecture perspective. The following subsections provide further commentary on the complexity, reliability, security, and performance of the options.

5.1 Complexity

The high-level functionality of the two options is broadly similar, with the clear exception of the SI communication in Option YHUB. Aside from that, both options include:

- The ability to generate, issue, and track references for each switch
- A mechanism to initiate a switch
- A mechanism to cancel a switch

5.1.1 Orchestration

While the two options provide similar functionality for orchestrating a switch, there is a difference in the approach taken for the role of the Hub.

In Option X, CPs interact with the Hub to obtain switch codes and update switch code status. The Hub is a central repository for the status of all switch codes, and in effect acts as a central control point. In the Option X approach, there is no need for the GCP and LCP to communicate directly with each other, and indeed, there is no support for this.

In Option Y, the Hub fulfils a different role in switch orchestration. In Option Y, the switch request is set up by the GCP and LCP exchanging messages with each other, via the Hub. The Hub in Option Y is responsible for issuing switching codes (switching order references) and tracking their status, but it does this from a reporting perspective, rather than having any active controlling role.

From a technical perspective, given the independence of the GCP and LCP interactions in Option X, it should be easier to isolate faults between the gaining and losing parties. However, this benefit needs to be balanced against the use of the supply chain in Option X for communication with the Hub.

Involving the supply chain adds additional parties and interfaces which adds complexity to Option X. The interdependency between systems is explored further in Section 5.2, below.

5.1.2 Message Transfer

In terms of complexity of the interfaces, both solutions propose similar interfaces and queuing systems for the various parties to interact with each other or the Hub.

The messaging between parties is asynchronous but needs to feel near real time. Assuming the use of a stable and well-functioning microservices architecture, it is unlikely that transient delays in message transfer would have a material impact on the process in either option.

5.1.3 Communication of Switching Implications

Option YHUB proposes that SI is communicated to the customer via the Hub.

The Hub will need to support communication of SI in durable form via email, text message (link to download), and letter. Additionally, for customers using an online channel, the Hub will display the SI in an inline frame on the GCP web page.

The above functions all rely on technology that has been available for many years and is well understood. As such, the technology does not introduce any specific or unusual risks.

Inline frames (iframes) are commonly used on the web to enable content from one web page to be displayed on another page. Iframe functionality is widely supported in modern browsers.

5.2 Reliability and Availability

5.2.1 Individual CP Systems

For both options, the availability of the individual CP systems and Hub will be determined by design and implementation choices, and the thoroughness of the testing. Some CPs may choose to design for high availability with partial or full redundancy; others may not.

These choices are expected to have a greater impact on the availability of individual CP systems than the decision between Option X and Y.

5.2.2 System Interdependency

Both options rely on having systems from different parties available at the same time. It is therefore useful to consider how the interdependencies between the systems will affect the overall availability.

As described in the Section 4 (Assessment of Business Risks), if individual systems are unavailable this will cause disruption in the switching process. Where this impacts a customer touchpoint, the affected customers may have to interrupt their switching journey and reattempt at a later time.

In terms of the customer journey, Option X includes two customer interactions that require multiple systems to be available (LCP to Hub, followed by GCP to Hub). As previously noted, the communication with the Hub is expected to be passed through the LCP and GCP supply chains.

Option Y has a single customer interaction for which the GCP, LCP, and Hub must all be available.

The system interdependencies to support these customer touchpoints are illustrated in the table below. As previously noted, retail CP systems are taken to include any TPI systems on which the CP is reliant.

Note also that back-end activities that are hidden from the customer have been ignored in this analysis. While these may also depend on multiple systems, it is assumed that CPs would be able to recover from temporary systems outages without impacting the customer experience.

Table 2. Systems Required to be available during Customer Touchpoints

Option	Customer Interaction	GCP	GCP Supply Chain	Hub	LCP Supply Chain	LCP
X	Get switching code			✓	✓	✓
	Validate Code	✓	✓	✓		
Y	Initiate Switch	✓		✓		✓

If the availability of the CP systems and Hub is assumed to be independent of the chosen option, then the following observations can be made:

1. The Hub occupies a central role in both Option X and Option Y. Investing in a high availability solution for the Hub will improve the overall availability of both options.
2. The overall availability of Option X is dependent on the length of the supply chain, i.e. the number of intermediate systems between the retail CPs and the Hub.
3. The availability of Option X gets worse as the length of the supply chains increase.
4. The degradation in the availability of Option X could be largely mitigated by investing in high availability solutions for the intermediary “Hub integration” services provided by members of the supply chains.
5. In the Option X scenario where the retail CPs connect directly to the Hub, the overall availability of Option X is equivalent to that of Option Y. (Each of the two customer interactions in Option X will individually have better availability, but the overall availability will be equivalent.)

5.2.3 Overall Availability

As discussed, the overall availability of the chosen solution will depend on the individual availability of the CP systems and Hub, and the interdependency between these systems within the solution. In this section, we provide a sense of scale to provide some context for the foregoing.

As it is not possible to estimate the availability of the individual systems – as this will depend on design and implementation choices – metrics from a recent cross-industry initiative in the banking industry have been used by way of illustration.[11]

Under the Open Banking initiative, banks have implemented Application Programming Interfaces (APIs) that enable bank account information to be securely shared with providers of financial services. The target availability for account provider APIs is 99.5% (i.e. less than 4 hours of downtime per month, including scheduled and unscheduled unavailability). Actual performance this year has been closer to 99.0% (i.e. around 7 hours of downtime per month).[12]

The following table illustrates what a similar level of availability for the individual systems in Options X and Y would translate into overall availability.[13] This analysis demonstrates the importance of the supply chain system availability in Option X.

Table 3. Illustration of the Contribution of Individual System Availability to Overall Availability

Option	GCP & LCP Availability	Supply Chain Availability	Overall Availability	Downtime per Month (hours)
X	99.0%	99.0%	96.9%	22.1
Y		-	97.9%	15.0
X	99.5%	99.5%	98.4%	11.5
Y		-	98.9%	7.9
X	99.5%	99.9%	98.8%	8.6
Y		-	98.9%	7.9

5.3 Security

From a systems architecture perspective there are well proven models for building the architectures presented. The adoption of architectures such as the TM Forum Microservices, and component architectures that detail secure approaches for microservices, should be considered. There are no material differences between Options X and Y.

From an operational perspective, Option X proposes a model where only APs and Voice over IP (VoIP) providers connect to the Hub, whereas Option Y proposes all retail CPs connect to the Hub.[14] Having fewer parties connecting to the Hub would reduce the chance that a connection becomes compromised. That said, this risk can be assumed to be low, provided that a secure method of connection is adopted in the design and access privileges are appropriately controlled.

[11] Open Banking figures are used for illustrative purposes and are not intended to set expectations for telecoms industry targets. Differences between the banking and telecoms industries may mean that these are not appropriate benchmarks for the telecoms industry.
[12] https://www.openbanking.org.uk/providers/account-providers/api-performance/
[13] The illustration assumes a GCP and LCP supply chains length of one, i.e. the LAP and GAP provide Hub integration for the LCP and GCP respectively. The availability of the Hub is assumed to be 99.9%.
[14] The OTA estimates there are currently 100 APs and 150 VoIP service providers versus around 600 retail CPs.

5.4 Performance

From a transaction volume perspective, the volumes proposed by the OTA (reproduced below) should not present any issues for a modern, scalable microservices architecture.

Table 4. Estimated switch transaction volumes from the OTA

Item	Current	Proposed build for
Switching (Inter and Intra AP combined)	~200k/month	300k/month
Retailer brands	~6.5k	10k (Note: the upper limit of the current RID format of 3 'alpha' characters is ~17.5k)
Retailer commercial entities	Estimate 600	1k upwards
Resellers / Wholesalers	Estimate 50	Suggest build for 100
Access Network Providers	Estimate 100	Suggest build for 200
VoIP services providers	Estimate 150	Suggest build for 300

Source: OTA

There are multiple examples of implementations of an architecture such as this. Many companies use microservices in a broadly similar way. An extreme example is Amazon's retail e-commerce offering, which is built using microservices that allow third party retailers to integrate into the platform and the distribution channels, and have the ability to scale on demand.

Another reference point is the previously mentioned Open Banking initiative. Since August 2020, providers in this distributed system have been processing over 500 million open API calls per month.[15] Applying a conservative planning estimate of 100 API calls per switch to the OTA's volume estimate (300k switches per month) the switching API volumes would be 30 million per month.

To account for growth in switching volumes, serverless architectures in cloud computing platforms that can flex with demand could be considered. If this is not possible, the solution should be designed with sufficient headroom to meet future growth, as indicated by the OTA. If there is uncertainty in future demand, more headroom could be considered, e.g. three to five times current volumes.

5.5 Upgradeability

It is possible that the switching requirements may change over time and therefore the ease of upgrading the solution is important. For both options, the extensibility of the microservices and messaging will be key to the design.

The use of an open message format such as JavaScript Object Notation (JSON) within the Hub and messaging is a good choice for system-to-system interfaces and will allow the message payload to be

[15] https://www.openbanking.org.uk/providers/account-providers/api-performance/

changed in the future, to cater for any additional needs. This is proposed by Option X[16] and could be used by either option.

In both options, CPs will need to conform to the interface specification for communications with the Hub (and other CPs, in the case of Option Y). Changes to this specification would need to be adopted by all affected participants. For Option X, it is possible that some changes may only affect those parties that directly interface with the Hub (i.e. APs, VoIP providers, and potentially larger CPs), but larger scale changes may also affect the interfaces between these parties and their downstream CPs.

6. Build and Deployment

There are two main considerations regarding the introduction of the new switching processes. First, whether a staged deployment is possible, and second, how to handle switching requests that are already in progress when the new processes are introduced.

On the first issue, Option X notes the attractiveness of a big bang approach, but also notes that there may be value in investigating a partial implementation to reduce deployment timescales. Specifically, Option X suggests that it may be possible for the four large retail CPs to cut over to the new switching processes before Openreach has completed its implementation. In this case, the retail CPs would directly connect to the Hub and continue to interface with Openreach as they do today.

Under this temporary arrangement, switches between the four large retail CPs (and any other CPs that were ready) would use the new switching process. Switches to and from other CPs would continue to use the existing process until they, and their supply chains, were ready to cut over.

Although not suggested within Option Y, a phased approach would also be theoretically possible for that solution.

The value of early, partial deployment needs to be weighed up against the potential for customer confusion. It is likely that some customers will not obtain switching codes when this would have been beneficial, and conversely, they may attempt to use codes with providers that do not yet support them. Ofcom may prefer a complete (“big bang”) cut over to simplify communication of the new process to customers.

For the second issue, a graceful approach will be required to complete switching and porting requests that are already in progress when the industry cuts over to the new switching processes. This is common to both Option X and Option Y. The best approach is likely to be a short period of parallel running to allow these in-flight requests to complete using the existing processes, with all new requests being handled by the new process.

Specific to Option Y, industry adoption of UPRNs must also be considered. Prior to going live with the Option Y switching processes, any CPs that do not already use UPRNs will need to map their existing customer address records to UPRNs. It would be ideal if retail CPs reconciled their UPRN records with those of their APs (and wholesalers and TPIs where relevant). Starting with a ‘clean’ consistent data set will reduce the likelihood of address matching failures.

[16] OTA

7. Reference Documents

The following documents were provided for this study.

- X Group_GPL Option X - Broadband and Voice Switching Proposal FINAL (OTA2) (1) _redacted
- X Group_Project EECC - Customer Journeys_FINAL
- X Group_Project EECC - Intra and Inter network switching flows_FINAL
- X Group_response to Ofcom Questions 200605
- Y Group_Answers to Ofcom Questions Issue 1
- Y Group_Option Y Revised Proposal-3rd Aug-Overall Summary _redacted
- Y Group_Revised proposal-3rd Aug-Process Flows
- Y Group_Revised proposal-3rd Aug-Process Steps
- OTA_Volumetrics Option X, Y and Hub issue V1.0
- OTA_EECC Switching Option - OTA2 assessment for Ofcom v1

8. Glossary

Term	Description
AP	Access Provider, the organisation that provides the network used by a CP
CP	Communications Provider, a provider of communications services
Customer	The account holder or other, duly authorised, end-user
GAP	Gaining Access Provider, the AP that operates the network used by the GCP to serve the customer
GCP	Gaining Communications Provider, the prospective CP that will hold the retail relationship with the customer at the end of the switching process
LAP	Losing Access Provider, the AP that operates the network used by the GCP to serve the customer
LCP	Losing Communications Provider, the CP that holds the retail relationship with the customer at the beginning of the switching process
OTA	Office of the Telecoms Adjudicator
PONR	Point of No Return, the point in an order lifecycle after which it cannot be changed or cancelled
SI	Switching Implications, information that CPs must provide to customers which explains the impact of switching including any charges and impacts on services
Switching Code	In Option X, a code that identifies the services to be switched
Switching Order Reference	In Option Y, a code that identifies a particular switch request
TPI	Third Party Integrator, an intermediary firm that provides business integration services to (typically small) CPs, including interfaces to wholesalers and/or APs
Wholesaler	A CP providing communication services to another CP

Cartesian is a specialist provider of consulting services and managed solutions to leaders in the global communications, technology, and digital media industries. For 30 years, we have advised clients worldwide in strategy development and assisted them in execution against their goals. Our unique portfolio of consulting services and managed solutions are tailored to the specific challenges faced by executives in these fast-moving industries. Combining strategic thinking, robust analytics, and practical experience, Cartesian delivers superior results.

For further information, please contact us at cartesian@cartesian.com