



UNIVERSITY OF KENTUCKY

Purchasing Division

Written Questions and Answers

Enterprise Integration Solution

RFP UK-2258-23

Closing Date: 07/20/2022

Today's Date: 06/17/2022

No.	Question	Answer
1	We noticed that there are a number approximately 90 source systems that are in use today and that there are approximately 230 endpoint systems and those appear to be utilizing the IBM messaging infrastructure. Is it UKHC's intention to replace the existing system with a new messaging platform? Or is the intent of this RFP being to add new messaging capacity and have parallel operation with the existing system?	UKHC's intention is to replace the existing system with a new integration platform. The number of systems were given for reference to show the current volume being supported today. Also, this shows the number of resources we will need to migrate from our existing system as part of this proposal effort.
2	Name of 90 sources and protocol used to connect to those through IBM products in the current setup	The majority of the current integrations are HL7 TCP/IP connections to a MLLP front end. There are 10 interfaces that are direct interfaces with X12 message formats over TCP/IP. Due to security concerns we will not name all 90 sources at this stage of the RFP (the selected vendor will have access to this information). The majority of interfaces come from Epic Bridges but we also support several EMRs from other facilities.
3	How much peak and average transactions per day in the current system	We have current transaction rates of: Peak per second: Inbound : 88 messages per second Outbound: 357 messages per second Total in and out: 445 messages per second Peak Daily: Inbound – 737,686 messages per day Outbound – 4,000,440 messages per day Total in and out – 4,738,126 messages per day Avg daily: Inbound : 563,790 messages per day Outbound : 3,202,954 messages per day Total : 3,766,744 messages per day

		Any solution should support typical year to year growth.
4	Is traffic continuous throughout the day or mainly in US business hours	The traffic is continuous throughout the day but between the hours of 7AM and 5PM there is more traffic. Peaks occur between 7AM – 9 AM, 1PM-2PM and there are some batches that run overnight but are not higher in volume than the day time traffic. Weekdays have more traffic than weekends.
5	Does current system supports Guaranteed delivery and is that going be a business need in the TO-BE system.	Yes, once the current system gets the message into a queue, we ensure Guaranteed delivery that is also first in, first out. If there is a failure before we get it into the queue, then we return a 'AR' Ack Retry message to the sender. In the event of a catastrophic failure we return an 'AE' Ack Error to the sender.
6	How many integrations/interface in the current system connecting to 90 sources	We have 223 outbound interfaces/integrations. 6 of those are FTP in nature and 10 are X12 passthrough flows. The rest are HL7 tcpip integrations that have filters and transforms unique to each endpoint.
7	Mapping : How many interfaces/integrations as very simple (less than 10 fields mapping), simple(10-20 fields mapping), medium (20-50 fields mapping), Complex(more than 50 fields mapping)	Of the 223 integrations approximately 25 are very simple, 110 simple, 75 medium, and 13 complex interfaces.
8	How many interfaces/ integrations are synchronous and asynchronous	We have 14 synchronous interfaces/integrations. All others are asynchronous.
9	How many UI/UX ?	We have a management UI which allows us to see the status of endpoints and servers. We also are using Splunk as a UI to see data about alerts, faults, messages in and out, etc. We have a UI that lets us search for messages that meet a criterion and display all the contents on one screen. Finally, we have IBM thick clients that allows us to see data about the queues and flows that are implemented.
10	is current system supports High Availability and Disaster recovery using alternate Data center	Current system uses VMWare ability to move virtual servers between hosts in two different data centers. However, there are some components of the solution that do not automatically recover in the other data center which is not ideal.
11	does current setup all on premise or it use some element of cloud and is hybrid	Current system is all on premise.
12	What are the different implementation/ frameworks to support system and daily maintenance	We currently have teams that total 9 people that develop interfaces and support test and production. We have all logging written to syslog on Linux and those logs are transmitted to Splunk. We have many dashboards in Splunk that allows us to see inbound and outbound messages by source/endpoint, alerts, and

		faults. Splunk is integrated with SCOM and ServiceNow to create tickets when issues are identified. We also have several tools that in place to allow us to monitor and administrate the system. We have a management ui that allows us to see status of all parts of the system and stop and start those components. All messages are stored in a filesystem on Linux for 90 days and we can search those messages and use those for replay and problem determination. We have thick clients that allow us to manage the IBM system. Finally, we have a UI that allows us to search for a criterion of messages and then display all messages that meet those criteria.
13	Is current system used only by internal users or is used by externally users/ partners over internet	The system itself is only managed by internal users. We do have partners that send and receive data but they have no access to the interface engine system itself.
14	Appendix 1 – Req. No. 10 Will you be deploying on-prem vs. public cloud vs. vendor hosted?	We are looking for recommendations and will consider all options as long as there is a Disaster Recovery strategy. Preference in this area will be given to solutions that support a hybrid model allowing for system to run in either a local data center failure or an internet/cloud failure.
15	7.2 – Alternate pricing What is the current daily inbound and outbound message volume going through your integration engine?	Please see the response to question 3.
16	How does data from end systems/interfaces are pushed to the current integration layer	The data is mostly pushed to the integration layer by the source systems with a TCP/IP connection to our MLLP listeners. These are all HL7. We do have some listeners on IBM App Connect Enterprise and they support the pass through of X12 data with an insurance clearing house. Finally, we have 6 FTP related interfaces that take a file that has been FTP'ed to the system and the system processes the file and creates HL7 messages to be sent to endpoints from the file. Some FTPs contain the whole HL7 message format and others are a csv and the system creates a HL7 message per row.
17	Please share architecture diagram for Integration landscape with source/target applications flow.	Due to our solution being custom built and its proprietary nature we will not be sharing an architecture diagram or list of sources/endpoints at this point in the process. Selected vendor(s) will be given access to this data at the appropriate time. The nature of the of interactions managed by our current engine has been described in other responses.
18	What is the volume of transactions/messages/requests (per second/minute) that the integration layer needs to support?	Please see the answer to question #3.

19	What are the number of environments that exists , DEV, QA etc	We have a development, 2 small scale QA environments, 1 “production like” test system and the production system.
20	How does your system provide the ability to replay messages that meet certain criteria either in a group or a single message? -- Is the expectation is to re-deliver messages based on a condition? Can you provide more information on this use case.	We typically have the following use cases in this space but sometimes we are asked for more complex cases that we have been able to support. At this time, we can replay for 90 days (with some endpoints configured for longer). 1) Replay all messages that were sent to an endpoint during a given time 2) Replay all messages that were sent to an endpoint during a time meeting a certain criterion (MRN, Name, etc.) 3) Replay messages sent to an endpoint while sending them through the transform again or sending exactly like they were sent the first time. (this allows for a resend after a transformation is changed). 4) Replay all messages into the system over a period of time for all sources or one individual source 5) Replay all messages sent to an endpoint during a time while changing the value of one field in all those messages.
21	How many people are on the integration team? We want to know how many developer licenses you will need in addition to how many people will need training on the new engine.	Our integration team is currently split over three teams that support the engine. Currently that is 11 people but we should plan for some small growth in the number of team members.
22	How much conversion assistance (if any) do you want? It appears you have 315 interfaces from Appendix2. Please suggest a % of those interfaces (if any) you would like us to convert for your team where the converted interfaces are then ready for application-level testing by your team.	With proper training we are willing to convert all these interfaces ourselves. At a minimum we would like to see 10-20 converted as a POC for us to then continue the process. Additional support would be required for difficult issues during this conversion. If a conversion tool is available obviously we would use that to automate as much as possible.
23	Can we receive a Word copy of the RFP? University of KY RFP - uk-2258-23 as well as the BAA - uk-2258-23appndx3.	No, sorry. We cannot give out Word copies.
24	Question 47 & 48 in Appendix 1 ask about skipping and removing messages. Can you please offer the use case implied in these questions?	The use cases are: 1) A message has a field value or issue that is forcing the endpoint system not to accept it. Ideally, we would have the option to change the value that is causing issue and resend. If not, we need to be able to skip/remove the message from the queue so subsequent message are not further delayed.

		2) Sometime we may find a number of messages were sent to a queue in error. We are able to search the queue for a string value and remove all messages in the queue that meet that string criteria.
25	Given the amount of detail asked for regarding the implementation and support plan, the July 4 holiday, and the expected responses from our questions not being received until after 6/14, can UKHC please consider a two week extension of the 7/6 due date?	Please see Addendum #1 posted today on the Purchasing website at https://purchasing.uky.edu/bid-and-proposal-opportunities
26	Can you please provide addition scope details about the "Technical Support" to be provided which is mentioned in Section 2.1?	Ideally, we are looking for training that allows our team to fully develop interfaces and support the system. Additional technical support would be required for difficult issues and prod support in a crisis where UKHC's resource cannot easily recover the system.
27	Is there a hard deadline for the new system to be up and running?	We do not have a hard deadline; however, we are looking for an aggressive implementation schedule. This would include training and an installed system where migrations could begin as soon as possible.
28	Are you considering any other cloud vendor besides Azure for the Enterprise Integration Solution? We understand that RFP 2262-23 indicates Azure as the cloud vendor for EDW Modernization Implementation.	We are open to other cloud vendors but would prefer to not have a solution in a vendor managed cloud.
29	Terminology mapping is not mentioned in the RFP requirements. Will you be keeping the current Terminology Mapping solution?	We will need to support Symedical for mapping. Today we support these with API calls during message transformation. We would need to continue this support or replace it with native Symedical connectors if available.
30	With respect to Master Data Management, will the current MDM solution be in place for the new Enterprise Integration solution to handle internal and external patient reconciliation?	Yes, we currently use IBM for patient and provider lookups. We are trying to shift the matching of external MRNs to internal MRNs to Epic and then send that match data to MDM for future lookups as well. So, when an external record comes to UKHC we can send it to Epic and Epic will patient match and then send out both the internal and external MRNs to MDM (through the engine) for storage of both. We use API calls to MDM to both store external data and do patient/provider lookups.
31	What is the target timeline (if any) to be converted off of your current integration engine to the new engine? Are there any specific dates we need to be aware of?	Please see the answer to question #27.
32	We have not submitted hard copies/USB copies of RFP because of COVID, since the beginning of the pandemic. With this	No, please submit your response as asked for in the RFP.

	in mind, is it possible to submit the final RFP response via email?	
33	Are you open for Fixed cost bid or hourly rate?	We are open to either model.
34	Are you open to automation technology (RPA)?	Yes, we are open to automation where it adds to the overall quality of the solution.
35	Are you open for Offshore resources?	Offshore resources have to adhere to governmental standards around VA data. Offshore resources are not able to view PHI data and therefore should be used in support of onshore resources.
36	Question regarding SWaM: Will using a SWaM Boomi certified subcontractor to deliver the services also meet the Univ SWaM requirements?	The University does not have any SWaM requirements but does have a 10% goal of addressable spend.
37	How many total endpoints does the University of Kentucky have? a. Can you specify what applications you would like integrated within the first 6 months to a year with your chosen iPaaS? b. Can you provide an entire list of your endpoints? c. How many endpoints are in scope to be integrated after the first year? d. Where are your applications hosted?	We currently have 223 total endpoints. We would desire all of these to be migrated to the new solution within one year. We will provide an entire list of endpoints to the selected vendor only. Currently, our application is hosted on prem in a UKHC data center that has VMWare failover capabilities to a secondary UKHC data center in a separate location from the primary data center.
38	How many test environments do you require?	We require one test environment that is production like but does not have to support the full volume of production so it can be smaller in scale. Even being smaller in scale it must have load balancers, multiple servers, etc. if production has the same, but if production has 6 servers for volume test could have 2 (but not one). Our test environment cannot contain PHI data. This test system would be connected to other test systems including Epic in the enterprise. It is referred to as our integrated test system for this reason. We also like to have at least one QA environment where we are not connected to real endpoints where we can test changes or issues with prod data. During the migration in the past we have created multiple QA environments so we can test multiple endpoints at one time to speed up the migration testing. Once the migration is

		complete, we usually cut back to one or two QA environments.
39	Is creating, exposing and being able to manage API's a requirement in your iPaaS Solution? a. If so, how many API calls/transactions per day are you looking to do? i.e. up to 100,000 per day, up to 1,000,000 per day...	Our current requirement is for the solution to be able to call APIs in support of transformations and data flows. We would like to be able to extend some functions via API as well if possible. We initially would expect low API volume in a day (100,000 or less) but would hope to see that grow and the industry moves from traditional TCPIP HL7 interfaces to more API interactions.
40	Do you have any high availability or fail-over requirements?	Absolutely. We need ability to stay up almost 100% of the time. We would hope to be able to failover to a different data center/cloud if required but would want failover available for all components of the system.
41	What processing model do you use? (ie. batch or real-time)	We are predominantly real-time but we have 5-10 jobs that may read a batch file and create messages to send out to an endpoint.
42	Can you share a few use cases including describing the flow of data for your most critical integrations that will be moved to the iPaaS solution?	Ultimately, we have a pattern-based solution at the moment so almost all data flows follow the exact same path and are just different in the filtering and transforms applied to each message. Here is the basic flow of our engine: 1) Source system sends in HL7 message over TCPIP to our MLLP java program that has a listener or listeners for each sending system. The MLLP uplifts the HL7 to xml and sends it on for continued processing. 2) We do processing on each message that stores required filtering values to the message. Some of the values are as simple as a field value but we also have some very complicated filtering criteria that is based on a calculation or smart processing of the value of multiple fields. For example, we may use the DOB to see if a message is for a patient under a certain age. Or we might say if a field has a value use it, but if it is blank then use a different field. 3) We then send the message to be sequenced to make sure it gets sent out in the order it came into the system. 4) Filtering is applied by a pub/sub model that looks at the values calculated in #2 and the messages are put in queues for each endpoint that meets the criteria. 5) We have endpoint flows that read the message from a queue and applies the required transforms/processing to the

		message. This can include DB lookups, API calls, and other intelligent processing. The majority of our flows also call a xslt transform which is where we try to keep the majority of our data manipulation. This flow then sends the message to the endpoint and handles retries every 20 seconds as necessary. 6) During this process we send a copy of the message to an archive that stores the messages for at least 90 days. We store the message exactly as we received it from the source, how the message looks before it is processed by an endpoint flow, and then exactly how it was sent to the endpoint (post translation and manipulation). Our system supports very complex transactions. While it is written specifically for UKHC, it is mature in many ways and the fact that it is “home grown” should not be mistaken for simplicity.
43	For your HL7 requirements, do you use EDI trading partners to exchange these documents? If so, how many trading partners does the University have total?	We do not have any official EDI trading partners. We do exchange data with a number of external partners including regional hospitals, transplant services, local clinics, and KHIE (Kentucky Health Information Exchange).
44	Questions in regards to services. For assessment of Time and Materials (T&M) estimates we follow a calculation of time based on level of effort and grading the level of difficulty (simple, moderate, complex*). Responses to following questions will guide input to the estimation equation:	
45	Which & how many end-points are involved?	Please see answer to question #6.
46	Which business objects need to integrate between these end-point systems?	We need to migrate the integrations between the systems. This includes database lookups, API calls, transformations, and custom code
47	What is the timeline for the implementation project?	Please see the answer to question #27.
48	Are there data lookups required in any of these integrations? If so, How Many?	Yes, we have approximately 15 interfaces that interact with Database tables either storing or reading data during real-time processing. We also have about 15 interfaces that have API calls to outside systems to get data for processing.

49	What are the process workflow requirements? a. Are there complicated steps to achieve the integration? b. Does the integration require customer alerts and error traps?	a) We have many complex integrations that require some custom code or complex xslt to perform. We also do some real-time queries where we take a HL7 message, query an API and return a HL7 message to the sender. These are synchronous transactions.
50	Does the customer administer and control these end-point systems or will other parties be involved?	While UKHC analysts have access to the vast majority if not all internal systems we support there may be some vendor involvement as well. In these cases, the UKHC resource will be the lead and the point of contact. We also have external sources and endpoints where we would need to have other parties involved for VPN setup and testing.
51	Can you share all key challenges you are facing with the current integration engine?	One of the main issues we are facing is the amount and quality of monitoring and auditing we have available. At current it is somewhat slow and cumbersome to search for all messages for a particular patient over a period of time. Also, we are slow in getting notified of actual issues in the system. Some of that is the delay in reporting to Splunk and some is the delay in Splunk getting it to ServiceNow. As the system is custom, we have found it is difficult to bring new resources on board quickly to learn development and operational support. It is very complex and takes a lot of in-depth knowledge to support the complex system issues. Finally , deployment of interfaces from test to prod is somewhat time and consuming and has enough manual components that it is possible to not deploy what was tested.
52	Can you provide a breakdown of the existing interfaces based on the following, and any additional relevant parameters you think would help us better estimate the migration? - Data flow (inbound vs outbound) - Type of source/target (internal/external, EHRs, clinical apps, payer systems, HIEs, etc.) - Format/feed type (HL7, CCDA, X12, FHIR, DICOM, flat files, DB links, etc.) - Complexity (S/M/C)? - Downstream integration (EDW, portals, real-time notifications, analytics, HIEs, public health reporting, etc.)	- We have approximately 90 inbound and 223 outbound interfaces - We have a strong mix of source and target systems. Epic is the main source and target but we have clinical systems including but not limited to lab, cardiology, radiology, and outside labs. We have billing systems, insurance clearing houses, HIE, and a variety of external interface engines that we receive from and send data to.

53	Can you share architecture diagram/details of the current home-grown integration engine?	Please see the answer to question #17.
54	Can you provide indicative daily data volumes, preferably broken down by source/feed?	Please see the answer to question #3.
55	Can you provide expected YoY growth in number of interfaces and data volumes?	This is extremely difficult as we are a growing organization and since we have now transitioned to Epic in 2021 we are looking for new opportunities for growth. I would assume at least a 10%-20% growth rate at this time.
56	Can you provide details of real-time / streaming data interfaces, if any?	We do not have any streaming data interfaces. We use real-time HL7 data transfer but do not stream.
57	Can you share any functional/technical gaps and/or planned enhancements (e.g. data quality checks, EMPI, etc.) that will need to be supported in the new solution?	We are looking for more monitoring and quality control capabilities. We are not aware of anything other items in response to this question but as a growing organization there are sure to be some.
58	Can you share the size and skillsets of the team that manages the current integration engine?	As mentioned previously we have approximately 11 team members currently supporting the system. The skills are technical in nature and all the resources can develop new interfaces and support the production system on a daily basis. All resources are very skilled in HL7 and the workings of an integration engine in a healthcare setting. One of the resources is focused on data conversions and would use the integration engine to assist with data loads to new and existing systems.
59	Do you have any preference of integration engines, maybe based on past experience, existing skillsets, investments, vendor relationships, etc.?	The purpose of this RFP to identify the integration engine that best fits the needs of the organization.
60	Do you have any hosting preference (on-premise/ private cloud / SaaS)?	Please see the answer to question #14.
61	Are you open to implementation and/or support/maintenance services being provided in a global delivery model (mix of onshore and offshore team)?	Please see the answer to question #35.
62	Are there any specific timelines/deadlines for completing the migration?	Please see the answer to question #27.
63	Are you planning to host a vendor briefing call for this RFP?	No there will not be a briefing call.
64	What additional technologies are used besides WebSphere for integrations?	We are currently using a suite of IBM middleware products and custom-built java applications. The engine solution itself includes IBM DataPower, IBM MQ, IBM App Connect Enterprise, and IBM WebSphere Service Registry and Repository. The logs are sent to Splunk for alerting and monitoring assistance.
65	Can you describe additional systems you plan to integrate with (ex. RDBMS, Salesforce, Workday, ServiceNow, SAP, etc).	We would love to interface directly with ServiceNow for alerting. We also use Xmatters to notify people and it would be nice to integrate with that. Otherwise we are currently integrated

		with our MDM solution, Symedical, and Microsoft SQL server.
66	Are there any preferences for UK's deployment model: 1. on-prem 2. cloud ecosystem w/in UK's control 3. vendor-supplied cloud environment 4. hybrid	Please see the answer to question #14.
67	Can you please describe the implementation timeline for the migration from WebSphere. Will it be a phased approach?	Please see the answer to question #27. We would likely do a phased approach where possible. The phase approach would typically be source driven as we would move a source to the new system and thus would have to move all the endpoints that get data from that source at the same time.
68	2.2 Background Information Can you please share the current high level integration architecture diagram to understand more on the integration landscape?	Please see the answer to question #17.
69	2.2 Background Information Please share the list of all Source and Target Systems to be integrated.	Please see the answer to question #2.
70	2.1 Intent and Scope We understand that installation /setup/customization of health care software like Epic / Cerner is out of scope for Vendor. Please confirm.	That is correct. The scope is specifically focused on replacing the current integration engine and the integrations it supports.
71	Can we know the timeline /schedule for this integration work?	Please see the answer to question #27.
72	Can we know whether any new integration platform (iPaaS) is identified and finalized, or vendor is expected to provide recommendations for new integration platform?	The purpose of this RFP is to identify the best fit integration platform for UKHC. Recommendation should be provided and included in any response.
73	Is the requirement to completely move out of IBM stack? Or UKHC would likely retain certain products like IBM MQ, DataPower?	We are open to all solutions provided.
74	Can you mention the current deployment model for the integration platform? I.e. On-prem or cloud or hybrid? Do you have any preference in deployment model for the new platform?	On-prem. We are open to all options but preference would be a hybrid model that would allow us to operate in the event of a local data center failure and also during an internet/cloud/vendor failure.
75	Appendix 2 Background Is it a correct understanding that you have both Rel-time and Batch integrations? [Appendix2: 300 interfaces]. If so, could you please share the count for each?	We have mostly real-time interfaces. There are about 5-10 batch interfaces and the rest of the current interfaces are real-time or query interfaces.
76	2.2 Background Information	The 230 endpoints mean different end systems and/or different interfaces to one end system

	Is it a correct understanding that ~230 endpoints means different end systems and not APIs endpoints? Are these part of the original 300 interfaces or separate?	(i.e. one interface for ADT data and one for Result data). These are a part of the original interfaces. The 300 number was given to allow room for growth.
77	Appendix 2 Is it a correct understanding that 35 ESQL functions need to be re-coded in new platform language, preferably like java or .Net?	Yes, however the key is that the equivalent function is provided. If the new solution can support the same functions without additional coding that would meet the requirement as well.
78	Do we need to support the current integration platform during the course of the migration?	No. The UKHC team that currently supports the integration engine will continue to support it until all interfaces have been migrated to the new solution.