

February 2020 Release Webinar

Question and Answer Session

Q: When will servers be upgraded to the new version?

A: The upgrades to hosted servers will start very soon. For those of you who are on your own servers, I believe the release is already available from our Downloads site. For the hosted servers, I think we just started upgrading some of the trial servers and the rest of the servers will start being upgraded in about two weeks or so.

Q: Is Google Translate available to all users?

A: The Google Translate integration is used to translate system text, such as table names and fields. You can make it available to administrators in a few different ways, depending on whether your administrators are expected to translate the entities they create as they go, or whether you have dedicated staff who review all the new translations at once. The integration does require a Google Cloud API key.

Q: Will the AI Engine be available for Professional Extended users?

A: Customers will need an Enterprise license to access the AI Capabilities in the AI Engine and an Enterprise Extended license to create custom AI with the open AI integration.

Q: Does AI only extract parties you already have or new parties as well?

A: The AI is essentially two pieces: there's what the AI model returns, where the AI will know where to look in the document to return the companies that are involved; and then there's the second piece, which is how you map that to the local fields in the table. In my example, since I knew it was companies we were already working with, we mapped it to our Existing Companies field. If you're not sure, if it could be new companies, then we would recommend mapping it to the New Company field in the contracts. Then from there, we could do an examination, try to match it and see if it's an existing company or not and update it as needed. So, it definitely can extract companies that are new, not just ones already stored in the system.

Q: For the Home Page charts and reports, will these changes happen to the existing graphs automatically or do we have to do something particular?

A: There was a bit of back and forth on whether we wanted to update the charts automatically, and I believe we decided to not update them automatically. After the upgrade, if you edit your existing charts and reports, there is a checkbox to say Use JavaScript Charts. There are a few chart types that we haven't quite added to the JavaScript charts yet, so there's a few less options, but we're continuing to work on those. And then with the JavaScript charts, once you enable them, you automatically get the mouseover effects and the ability to click on the legend to update what's displaying in the chart, as I showed. The drill-down pieces have to be added to each chart, so in each chart, there is a new tab called Drill-Down where you can specify what you want it to do when you drill down into the other layers. So you kind of specify what other fields it's going to look at when you drill down.

Q: Can we see what's in the source document to see how well the AI pulls the clauses info?

A: Here's one of the documents, the one I used the metadata extraction on. You can see YM Biosciences, Allphase, and our start and end date. If we scroll down further, you can see the signer that it pulled out.

Master Services Agreement For YM Biosciences

This Master Services Agreement ("Agreement") is entered into as of May 17th, 2007 (the "Effective Date") by and between **YM Biosciences Inc.**, a Canadian corporation with an office at 5045 Orbitor Drive, Bldg. 11, Suite 400, Mississauga, Ontario L4W 4Y4 (collectively with its Affiliates hereinafter the "YMB") and **Allphase Clinical Research Services Inc.**, having offices at 19 Camelot Drive, Ottawa, Ontario, K2G 5W6 (collectively with its Affiliates hereinafter "ALLPHASE"). YMB and ALLPHASE are referred to singly as "Party" and jointly as "Parties" throughout this Agreement.

WITNESSETH

By:	By:
Name: Jess Smith	Name: David Allan
Title: President	Title: Chairman and CEO
Date: 22 MAY 2007	Date: 16 MAY 2007

Q: How do you teach the AI to differentiate between document types and how to extract metadata from a document?

A: The short version is that you set up the new Machine Learning action type to perform training on the machine learning model. You can also configure that action type to ask the AI to analyze the data instead of using it to train. When you use the action to perform training, the action sends the data to the model to train on it. However, before you can train any models, you need to prepare training data to feed it, which is a more involved process because someone needs to look through a bunch of documents and tag them. For example, to train a model in metadata extraction, you need to provide a set of documents that have the metadata identified and tagged, so the model can learn to recognize the

metadata on its own. We've had our teams working on this for a while using a lot of our sample documents and common document types out there.

Q: How do we get out-of-the-box AI and machine learning functionalities, and any other functionalities that weren't available when we first implemented our system?

A: For customers with the appropriate license, the actual AI functionality is available as soon as your server is upgraded; it's just the configuration of the system that doesn't automatically update. So the AI functionality, plugging in the AI models and using them in the new action types, will be available to everyone with an Enterprise license. However, with those pieces, some configuration is still necessary. There were a few configuration differences in the project I demonstrated from, for example a table called Clause Modifications. When I did the clause extraction, all of the clauses that were pulled out were created as records on a new table called Clause Modifications. That new table would not be added to existing projects automatically. We're working on creating an entity set with all the pieces I demonstrated so we can roll those out to people as they want them. So, the new table wouldn't necessarily be immediately available, but for example the metadata extraction would be, because the metadata extraction is actually just within the one Contract table—assuming everyone has a Contract table. It doesn't use any new tables; you basically just have to apply the mapping. As I said, the AI model is already going to be there, you just need to create an action that applies the mapping to the local fields.

Q: What's the expected average time effort investment for training the AI models?

A: It's difficult to give an expected time effort for this because there are a lot of variables. In a general sense, once we have the data, sending it to the AI to learn from it is relatively quick. The longest part of training is creating the training data set, which involves going through contracts and tagging the items you want the AI to identify: telling it here's where the company name is, here's where the signer is, this is what a clause looks like. Having good, high-quality training data is essential for getting good analysis from AI. Currently, it takes a couple of hours per document to tag and prepare the file and run it through the model for training, but it takes a couple hundred documents to get a sufficient body of data for the AI to work effectively. That number is coming down, but it's still a significant effort to train an AI.

Q: Can we see how a change in Microsoft Dynamics would reflect in Agiloft?

A: You can send information both ways. This demo is currently set up so Agiloft is the system of record, so it pushes its changes to Microsoft Dynamics, but you can set it up both ways.

Q: How do users' abilities to edit records and fields via Alexa reflect or build upon their abilities to edit records and fields directly through point and click?

A: The user's ability to edit records and fields through Alexa should be the same as they are in the interface. It does use your particular user account, so Alexa was acting as me, and if I ask Alexa to edit a field I don't have permission to edit, she'll tell me I can't do that.

Q: Does it require a Microsoft Dynamics license to integrate it?

A: Yes. We aren't providing a Microsoft Dynamics account. This is integrating with an existing account.

Q: Will there be an AI specialist rep assigned to every client? What will AI support look like?

A: We have a special AI team that has led the process and internal training to bring teams up to date, but otherwise they will be available as we begin to answer more advanced questions.

Q: Is the AI supervised or unsupervised?

A: Our Named Entity Recognition (NER) and Clause Extraction (CE) models are Supervised. We support an open integration with SageMaker to allow use of any of their models, and some SageMaker algorithms the customer can train are Supervised, some are Unsupervised. Agiloft is investigating the use of Unsupervised models, but these are not included OOTB in the new release.

Q: Does the AI extraction work with Adobe PDF or other non-Word files?

A: It works with Adobe PDFs and images. There's definitely an OCR portion to it to take non-text files and process them so they can be extracted. I haven't tried other formats like Excel, but we've definitely used PDFs in addition to Word files. Keep in mind that the readability and quality of those files is a factor, so PDFs with lots of handwriting on them or files with lots of photocopy noise are going to have less reliable output.

Q: Will the system be able to auto-assign follow-up dates based on timeline clauses?

A: There's a couple steps that would be involved. First, we'd obviously need clause extraction to pull out the timeline clause. Then, you'd need some configuration changes around that. You'd also need to pull out the date, so it'd be a combination of methods. It's within the realm of doing with a complicated setup.

Q: Is provided data used to train the algorithm and is it anonymized?

A: I believe there are going to be options to use the grouped data model we've been working on and feeding examples from online, or you can use your own fresh model that has never seen anyone else's

data. Any data that is used in the shared version will of course be anonymized so no one else can access it, but it's up to you if you want to use that shared data training or not.

Q: Do you have full text search over the whole system and its contents?

A: We do. We usually use full text search on a specific table, but there is a global search that looks through all tables.

Q: Do you have a certain hierarchy for contract master agreement and have all its supporting documentation and amendments linked to it somehow?

A: Generally, for contracts we have a parent contract and then any amendments or sub-contracts are linked to it and appear underneath it. We generally have them as separate records; there's actually a diagram on the Related Records tab so you can see how all the contracts relate to each other. And, typically, there is a master service agreement with a number of different contracts underneath that. Our Sales or demonstration team could show you more on that.

Q: Do you have a phone app?

A: No, we do not have an app that you need to purchase, but you can easily access your Agiloft system through our mobile interface. You have to configure it with where your project is located and how to log in. That gives access to our mobile interface. You don't have access to admin setup screens, but you have access to pretty much everything else.

Q: Can the risk analysis flow up the contract record to the company?

A: The demonstration of risk analysis I had done started at clauses and pulled up into the contract. Of course, the contract is for a particular company, so we could pull that information up into the Company record. Depending on how you want to configure it, you could either do the average across contracts for the company, or maybe the risk for the master service agreement. Or, if you expect to have just one contract for each company, it could show you those numbers. To start with, in the company record, there's already a display of the contracts, so I would just show those risk numbers alongside each contract and use row coloring to flag high-risk contracts.

Q: Do you integrate with Coupa or DocuSign?

A: Yes, we have a built-in integration with DocuSign and Adobe Sign for e-signature. For Coupa, we don't have a built-in integration, but we have done script-based integrations with Coupa for several customers already. The difference is that you won't find Coupa in our setup screens.

Q: Can external users go into the system and see agreements?

A: That's totally up to you and how you design your system. We've built it differently for customers with different needs. We have some customers that are public entities, such as local governments, and they're required that every contract they have must be available for the public to view. For that, we create a Summaries table that serves as a public view of the contracts, and that's what external users can come in and see. Others are of course not public entities, and they don't want anyone to see their contracts, and they lock their contracts down very securely even for internal users.

Q: Can you integrate with SAP HANA database management system?

A: We can integrate with it using the APIs, but there is no built-in integration available.

Q: Can we send information both ways between Agiloft and ServiceNow or SharePoint?

A: We don't have a built-in integration with ServiceNow, but we do have one with Jira. For individual customers, we built one using the API. For SharePoint, I know we've read information from SharePoint, but we haven't sent information to SharePoint. Neither ServiceNow nor SharePoint offers a built-in integration.

Q: How does the versioning and tracking of edits to agreements work?

A: In the general sense, when someone edits one of the documents using direct editing, every time you make a change, it's basically saving a point version back in the cloud. Once you release it, it saves a new version check-in. In my demonstration, we added a rule that says when there is a new version, run a conversion to create a new record. In terms of tracking edits, this is simply whether you have Track Changes turned on. If you're concerned about it, we recommend turning Track Changes on and locking it.

Q: Is Alexa available while using the browser on an Android device?

A: I was using Alexa through my Android device. I don't know if it ties into the browser directly. It depends on the phone settings. Depending on your settings, you don't necessarily need to use the browser; Alexa communicates on its own.

Q: Is there guidance on developing the risk formulas?

A: Yes. Our team is getting up to speed on the new release, and then they will be prepared to answer questions on this. We have some ideas we're using, and we're also getting up to date on best practices.

Q: If Alexa knows it's you, isn't that a personal account, not a work account?

A: Basically, there's an Alexa skill and then a key was generated that I put into my Alexa account. If you have a work phone, it would be on there. I do understand the security concerns—phones tend to be less secure than workstations. Remember you don't have to use Alexa. In order to start using it, you have to set it up in your system and generate keys. For more information on Alexa, please read the [wiki article](#).

Q: Does the system support approvals and rejections via Outlook?

A: Yes, you can embed links in emails so you can offer a single-click Approve link. We also include a link to review the record in the system, where you can leave rejection notes if you want.

Q: Is risk assessment an out-of-the-box feature?

A: Not really. We're building it into our new project base so people who start new projects will have our default configuration for risk assessment, but the extra fields and calculations aren't automatically added to existing projects. The assessments can actually be added by your implementation specialist, and we will probably have an entity set for those to roll this out more easily in the future. In general, we are always going to need feedback from the customer on what constitutes high or low risk. Even an out-of-the-box standard would need to be tweaked for each customer's needs.

Q: Does the AI learn over time and correctly tag items previously marked as not confirmed?

A: Yes, the AI is always learning whether you're training it or just getting responses from it.

Q: Do you have a plug-in for Microsoft Word, and does Word connect to the clause library?

A: I demonstrated our direct edit for Microsoft Word; that's currently our limit for Word integration. In a few cases, we've enhanced print templates for customers so edits to print templates can be put back in the contract. We're working on a further integration where from Word, you can integrate with the clause library and feed information back and forth.

Q: Does the system have a side-by-side comparison of agreement versions?

A: In the Attachment records, we can compare versions of a document and create a redlined version of it. You can compare just the previous version or select an earlier version if you want. You need to enable our Word API to do this.

Q: Does the system auto-fill data via a Salesforce integration?

A: If you have an integration with another system like Salesforce, then yes, we fill based on the information coming in from that other system. Depending on where the data is coming from, we would automatically fill it in. From the end user portal, there's also often quick links to fill in standard language, like a link in the request portal for an NDA would fill it in for them.

Q: Is the data behind the dashboards available to export to Excel or PDF?

A: Yes definitely. The dashboards and reports are pulled from data in the different tables. The reports themselves, you can generate the chart as an HTML version and as text or Excel versions of the output. Additionally, the table data the reports are based on can be exported to text, Excel, or even XML or JSON.

Q: Can the GUI colors be changed?

A: Yes of course. The look and feel can be changed, the logo can be changed. For the Alexa demo, I had the Alexa logo set up. If you go into Setup Look and Feel, you can set up different themes for different teams of users. You can change the colors, logo, even the font, and all of that.