Male Speaker 1: All right hi folks welcome to IBM Cognos enterprise planning solutions powered by IBM Cognos 8.4. So this is our introduction to all of the new and great features in Cognos planning 8.4 that so many of you are upgrading too today. I am here with my colleague Martin Overton and we are going to take you through well some of the, the flashy things that you may all be already be aware of like the new rich client as well as some of the more detailed technical pieces job server clustering changes and all the fun stuff. So I am going to get you started and we are going to walk through the new contribution experience. So I will show you new rich client and all the things that are available there and Martin is then going to take you through the installation and provisioning of that client because we had changed technologies and some of the other new features you might not be aware of, as well as go through some performance highlights and conformance changes that we have made for this release so you be ready to upgrade to that. So without further due I want to jump right into a demo of the new reach client not bothered looking at slides or anything like that let's just get in here and take a look at that new rich client. This is new rich client and you can see here that we have persisted the same sort of tab navigation that you might be familiar within the active X client I can navigate from my asset purchases to my depreciation all that sort of thing. So we still got a tab based navigation, but what you will find is this does look a lot more like analysis studio and in fact that's one of our three key themes for this particular release, so the one that Martin is going to cover for you or one of our first themes was to replace the technology underneath. We know that active X isn't a technology that a lot of our T-departments are comfortable with. So we wanted to strip out that active X technology and do this rich client in an eclipse Java platform but I will let Martin talk a about that stuff and I will tell you about the other thing. So the two are really to get in more consistent look and feel with analysis studio and Cognos 8 so that users don't feel like they have drifted away from something like they are in a whole new place, this has to feel like what they have already been accustomed to. So you don't have to worry about training people so much. We get the similar experience as we would get with Cognos 8, so for example the drop zones and dragging and dropping dimensions around. I can just pick that guy and put him wherever I like to get the view that I want. That's been one of the key themes to make sure you know we have this dimension bar and we can navigate through the tree based on the context in that dimension bar. So I can easily flip versions and that sort of thing as well as create any sort of nested view of that data that I want. So go through some of the other more detailed things in the third theme really was to come up with a lot of the features our customers I have asked for, for years. So for example a classic one is the ability to enter very long descriptions, to describe or give a justification for something that I want to do and now we have word wrap to support those very long text entries and as well we can also easily choose dates when we want to purchase those assets or do anything else, keep it simple for the end user and make sure that they are not putting dates in the wrong format and that sort of thing. So that's one of the themes but again sort of back to the navigation and exploration and analysis you know we are also able to focus in on exactly what we want to see and work with in this particular plan. Key to this plan I am working at asset purchase plan right now is you know the big ticket items that I want to buy this year. So I am going to sort on those to make sure that I see the big ones first. So we can easily sort that data and get a view along that to the big ticket items, the higher salaries and my salary plan all that sort of thing or we can also

sort by alphabetic or ascending or descending for text and things like that. So I could easily look at text in a particular order and one of the other things that we have heard from customers you know finding a particular item in a dimension, a long dimension we will show you that in my employee plan. Now if I want to look for a particular employee we will talk about tree view in a second if I want to look for a particular employee it helps if I sort those names to quickly find the employee who I want to plan for and give them a big race. So we can do that but to step back for one second to what we just blinked past there is the ability to see that data in a tree view. So let me look at this version here and I could now see a tree view of this dimension because there is a hierarchy in that dimension. So I can look through that and I can find a particular employee and expand based on where he is in a particular call center or department or what have you, and any dimension that has simple aggregation like this will be represented as a tree, but as you saw me do before if I want to change that tree to just a flat list that I want to look at I can do that as well. So we can do things like that we can expand and collapse. So we put that back lets' see let's go to our revenue plan for a second and I will show you another tree and try to highlight why I think trees are really important to the ability expand and collapse these hierarchies. So for example I have got drills here that I am selling this year and I think I am going to sell about 3000 drills based on a particular profile of what I have sold before. Now it maybe that we develop a plan with this level of detail in almost down to this skew level of products and that sort of thing but that's not actually the level that we want to work at. When I enter my data, I only care about how many drills overall I am selling this year. I am going to sell 4000 drills this year and if someone wants to see the detail of how that breaks out at the lower level I can always expand that and see the impact on the children. So we can find and you choose which level of detail you want to work at. I can customize this view in a host of ways. Some of the things I want to do with it though are typically to get rid of this stuff that I don't need. So I don't need January through what April, we are still in May so we work on that and hide these guys and now all of those columns are right on my way and I can continue to work on that plan. I can also unhide those guys and get back to the detail you know next year when I need to see what I did for the first few months. So that sort of thing we can also if we want a different way to sort of navigate around and get the view you want you can also freeze paints to that particular item and then I can scroll across and I maintain my view that the same way you would expect to happen in excel. Some great new features for customizing my view and all that sort of thing, how I want to see my data. One of the ways that I most often want to see my data is and coming up with the plan like this before I submit this guy on, one of the things I want to do is always have my data entry focused on the target you know I and the price I want to see my revenue plan that's going to be my target while I work with the rest of my plan or more importantly maybe we will go to even the more bottom line I want to see my income statement or put that guy over here. I want to keep my own income statement while I do my employee plan. So if I am going to give really here are raised to a grade F, I want to instantly see the impact of that on my income statement and know exactly how much that's going to cost me and how it affects my bottom line. So we can create all kinds of views here. I can drag these tabs off and put them anywhere I like if I want my expenses down here, just drag and drop back to other and as much screen real estate as you have to look at all these views you can create as many as you like. Its about this time everybody

always asks me what happens when I close this cube and I come back, I have spent some time nesting dimensions the way I want to see them reordering my tabs coming up with these views where I have got some re-cubes here and input cubes here all that sort of thing. When I come up with that view what happens when I close it? So let's close it and see and its at this point particularly when we close not when we save, its when we close that that view get saved for us and the next time I open up this client I am going to get the exact same view right down to the very cell I was working on when I closed it. So very easy for me to get back to where I was. Now see I am back in January 3 on this particular cell with this item highlighted over here and this one down here, anything I have left here is going to be there for me when I open it up again and the way that's going to work force there is three different ways we can deal with persistence. So the first one the one that we put in by default is that this wall must be saved to the server. So if for some reason you happened to move from one machine to another that's going to persist on the server. So when you open the application up you are going to see this view no matter which computer you have used, the other option if you like is to do local persistence. So we will save all your view settings here locally and we can do that in the third and final one is that we don't save them at all if the administrators really want to lock down these views so that people aren't creating you know lots of different ones we will still be able to create them but not save them, so we always close and reopen to get the administrator's view. So we can do that as well and on the note of you know these views that users can define and how complex they can get, it is entirely possible that a user might build by accident even the clicking and dragging around a view that they don't like, they don't want and maybe they don't know how to get back from what they have built, that's where they are going to call you and you are going to tell them to reset those views just from the view menu and if you don't know which they have, they have done if they have done some nesting or perhaps they dragged some tabs on just tell them to reset both that's going to take you right back to the initial state when the administrator defined how this application should look. That is the bulk of the new features that we have added in 8.4 and now I am going to let Martin take you through all of the fun server side changes that we have made and the provisioning of this new rich client.

Male speaker 2: Thanks Trevor. So we couldn't have, we couldn't delivered all those great feature that Trevor just shown here without making some technology investment here. As Trevor hinted what we have done is we switched our technology base from being the old Microsoft Active X standard that we have used in previous versions and we have now standardized on a rich client built in Eclipse and Java and that means that we have also been able to satisfy some deeds and be able to remove some of the IT objections that I know has made it tough for you guys rolling out the applications in your user base. So to say that no active X controls there is no common registration and it does mean that we are able to deploy this client now to users less standard Windows users. We are looking to fairly lot of ground environment you know longer need to be a local administrator to install the software. You still got two options you can even push the software out or you can pull our client and provisioning automatically but in the case of pushing it out we have standardized the installs to the Windows and installer technology. So we provide MSI packages now and you no longer have to create that in your organization. Okay so when we install the client there are two things that we need to put

down on the system really. The first is what we call a rich client platform. This is, this is a thing that includes all the plumbing, all the frameworks. So a Java room time environmental shift and also the eclipse framework that we are building on and the third thing is highlighted on the slide A as a provisioning agent. This is a piece of software that IBM Cognos have written which controls the download and the install of clients on top of this platform and the first client that we have written for this platform is the contribution client and so which Trevor has just taken you through. Now you may have noticed that I said the first client that's because you know its entirely reasonable in the future we will deliver with the client on top of this platform. I can't talk about you know what we will or won't do in the future but this is very definitely our new standard technology. Okay so in terms of getting that client out to the users and I have mentioned that there are two approaches and the first one being the ability for IT centrally to push that client out to the end users. So the rich client platform and that contributor client actually shift with the products as to MSI's, one for the rich client and one for contributor and then this gives IT the ability to use their system management tools of choice and push these out very quickly. Now I know in the past what the old installs people were tended to be created MSI packages themselves and they are no longer have to do those steps. If you want to install this client quickly we do still provide a client CD that you can just put in the CD drive or on the network I think you will get everything in that in one go with the UR based install. The other approach to deploying the client as I mentioned is the pull scenario. Now this is the one that we very much encourage you to try and standardize on. It has a number of advantages as we will see in a moment. I just want you to take you through the mechanisms of how this actually works. So when a user first goes into the Cognos 8 portal and they navigate to a Cognos planning package they see that the tree view. Now you have seen Trevor launch this client from this tree view but it's a very different from previous versions. So this tree view is also being rewritten and at this stage the client has now a software installed on their machine at all this is a zero foot print tree view. It's not until they actually go and click and open a note that they pull that Cognos RCP to MSI file down from the server. They will pull this down through the browser, they will choose to run it again I remind you that this will work even as a standard Windows user. Now it will give them the rich client platform. From there on this provisioning ageing component takes over and that will actually go to the server to the update site that you can see on the Cognos 8 gateway on the right there and that will pull down all the Java files and plug it to the client and you will see the progress as these are pulled down and that's how we get the contribute to client on top of that platform and the reason why I say that this is very much the approach that we want you to take if possible is that it does give us the big advantage that in future if you want to take fixed packs on maintenance really says future upgrades we can just simply install the updater kit server side on the gateway and on the application servers and that will drop the new files in this update site. Then the next time the client connects to the server they will just bring down the change files. So instead of saying you know downloading 1 of 89 as we see here on an initial pull, when they just want to take the updates quickly if this two client files have changed they will just pull two client files down. So if you have got you know a thousand desktops deployed over this is a much quicker way to be able to take those, take those later versions and be able to you know take your clients on the upgrade journey with you very, very quickly. The options that control this provision that say

whether you are going to allow the automatic downloads and updates they are in the admin console, we are in the system setting web client setting screen you can choose to allow the initial install or you can just choose to allow updates because one of the things which we find is popular with customers is they would like to push out the initial installs and then still allow the users to be able to take updates in the future and by the way the classic client is not going away, so all that was giving this new client technology if you also upgrade to 8.4 and then face the implementation of the rich client perhaps you think that you know the user's need to be trained or you need to provide new instructions on use of the rich client you can switch the classic client back on, on an application by application basis, again in the contributor administration console. So 8.4 the new client is obviously the big feature but there are lots of other things which you have been put in that which you are going to significantly out here as well and one of... I would like to add a few of them now and the first of them is some announcement to the deployment feature. So you remember we introduced deployments, inversion 8.2 but we made a real investment in this feature because there has been so popular among the customer base and some of the changes that made well how you use this a lot more on speed or your ability to promote change between deaf test and production environments. We have also switched to deployment as being preferred upgrade mechanism, so now if you are upgraded to 8.4 from 8.2 or 8.3, you can simply impart a deployment archive that you exploited in the previous version and we will upgrade all that content at the same time that we will be bring in. So let's switch to a demo and let's see this deployment input in action and I will highlight one or two of the new features for you. So I am bringing in this deployment package and straight away you can see that this was exported from version 8.3 of the product. So naturally as I bring this and I am going to be doing an upgrade of the content and you will see how quick and simple is to upgrade from 8,2 or 8.3 to the 8.4 version of the product. Okay one of the advantages of upgrade and why deployments is that you don't need to bring in all the content from your deployment archive. Over the time we kind of get clue to details in our source environment you know maybe there is some test application that you created, some out link, some out macros you might not want to bring everything, everything over and being selective about this content which you bring in or which you upgrade is a big benefit. So I am not going to bring in these *[Inaudible] libraries[Phonetic]*. I have already got those but I am going to bring in a couple of applications a link which moves data between those two applications on a macro which rooms that link. I have the ability to reset security as I bring deployments in and it's quite common for people in test environments to *[Inaudible]* users and not of the full production names base available and in this case I am just going to change one of the users from administrator to my Martin user. Okay there we go, so any notes in the applications that we previously marked to administrator now that I have brought them into my new environment they will be mapped to Martin. One of the new features that I want to highlight is this *Maple[Phonetic]* functionality. This gives me the ability to set properties on objects that I am bringing in from just a single place. So I no longer have to go for all the applications and rename them if that's what I want to do. This can also be very useful if you are reaching a period well over and for example you want to copy the *FYO8 [phonetic]* applications to *FYO9* this can be a very quick way to achieve the application copy for the new year. I can also choose a job server, or job service cluster to assign these applications to and that's going to give you some savings

in the target environment when we brought this content in. So that's all out I am just going to press finish now and this tells me that I can go and monitor the status of the deployments in the monitoring console. So I have got a deployment tab here and I can see the progress of the content coming in here. So what's happening now is its actually bringing those applications and its creating a new database containers for those applications and then its going to bring in the link on the macro and its also going to run a job to validate that link and see is they are running the validate link to job now, because I have already named the content that I brought it in that link will automatically be removed and I am sorry be pointing to my application names that prefix with the UP and so I have got you know no set of that to do whatsoever this content is now just ready to run. So that's finished and if I go to tools and refresh console, I can see my two new applications here. I will also see the link which is mount source application up_go *[Inaudible]* and target applications up_go *[Inaudible]* on the macros also to rule not new link as well. So just to close out this demo I am just going to GTP one of these applications. So I will create my Cognos package in the website and then its ready for users to go in and access. There we go, so that's complete to none if I drop into the web portal, log in as a matching user around *mapped[Phonetic]* content too during the deployment import and now I can see my new application. All ready to go, so I think you will agree that that is a pretty quick transition in the but old days you know it's a migrate content between applications like that until grade would have required database back ups, it would require transfers of *LAE [phonetic]* files to bring the security content across you know but this is really quick and in just two or three minutes down on here you see that this is up in running and ready for your users to work within the new environment. Okay so one or two other things that we have done this allocations within administration links this has been a very common request although the last few years from customers we were able to put it in this version and I know its going to be very popular feature. So now when you build an administration link it come to allocate yourselves and target allocations you can actually choose an allocation table from analyst to maintain those link mappings. That means that you can control these mappings externally via macros and keep this link in sink and keep this link valid as your dimension content changes in the application and we provide some new macros for synchronizing these allocation tables and also for validating the link is in a good state and this validation exercise also has a macro associated with it and that can be a great thing to put in as part of your overnight process, so that you can make sure that the link is going to work before it actually runs and you don't fall in a state of understanding that the link is broken when you actually come to the power failure. This is going to be a very popular for you too I am sure. So the question everybody asked performance and scalability, "How is the new version going to perform?" Well there is some good news here so the rich client that Trevor was showing earlier I mean this is a very different technology so what we had before. So its not easy to do an exact like for like comparison because there is a lot more going on into the covers and clips and that you have a room time environment they have a... they you know that takes a little time to initialize but it is a fixed time period. Previously if your application perhaps open in three or four seconds you know you may seem some overhead here you know this may go up to 8 or 9 seconds perhaps, but conversely to that if the application is large and was slow in the past particularly if it had a lot of items on rows or columns we could see quite a significant open time saving. So you know taking

an average across applications that people allow it to have you know I would say its fairly even but I don't want you to bear in mind that you know the technology is different you know and it can't, it can't vary by application but the real good news is that we have done a lot of work on the calculation engine and its release and so once you got that note open things that generally require a lot quicker. So large calculations, lot of brave bucks saving the data back to the server is across the board quite a bit faster. We are estimating in the labs that this improvement is around 20% and we also get that's the benefit of that improvement in the job system. So jobs that reconcile you know we are going to see an increase in performance there. We have also tuned the jobs themselves that you should see denied. *Locks[Phonetic]* will spend more time working and less time waiting. So you get better utilization on your job service and also in publish we have made some significant changes to the way it publish work, so that the main thing that we have done in publish is we have introduced some options to allow you to manage the indexing. So let me take in so as you know planning stores the data in a *[Inaudible]* format which is optimized for aggregations and calculations when you want to report on that data, we need to take that data out of relational forma. So what we do is we struggle out data to text files and then we book log.data into the database that's what publish does. Now in order to improve reporting performance we do lot of indexes in those tables in the published database. Through some testing in the labs we were looking at how we could improve the performance of loading that data into the database and what we found is that by removing those indexes the performance time was a lot better. So the restriction to publish jobs so the start of the publish job now we actually drop those indexes and then we load that data using the bull load tool optimizations and then we recreate those indexes at the end. So as far as reporting goes the data in the same structure the indexes are there everything is left untouched but we are able to observe some fairly significant performance improvement by doing in this way. The fastest way to publish directories if you don't need indexes in those published tables, so if you are *e-tailing [phonetic]* that data straight to a data warehouse and not reporting off it directly the *bestest[Phonetic]* option in this case might be to disable index in all together and the option is there to do that. So for a combination of these changes I have been talking to customers and you know there has been some significant savings reported on overnight performances you know saving tons or minutes or even hours of some of the longest processes is not uncommon start or the combination of data improvement. So this is an example for Oracle but what we found is because the index wherein on those tables when we were bulk loading the data we were able to run with the optimized parameters for the bulk load tools in this case *for an article[Phonetic]* direct equals true and parallel equals true and we have also done the same optimization equivalence for DB too and sequel server by the way and this graph illustrates quite nicely some of the results that we have seen. This is showing us the number of processes of threats that were way to commit data. On the left hand side of the graph way you can see that the graph is quite high this is when we were doing the publish with the indexes in place on the target tables in the publish data store. We are using 16 processors there in the labs and so the orange line there tells you that's you know really between 13, 14, 15 thread sessions were waiting to commit data during that publish operation. By dropping those indexes and using those tuning parameters to the build load tool we were able to load that data in parallel, no table locks were being taken due to the index and that data could be pushed in that very quickly. So they run

inside of the graph shows you that no commits were taking place. That's what has enabled us to increase the published performance and the figure you see that 50% is not an uncommon story from our testing in the labs and I got to say that some customers have come up to us and validated those results in their environments as well. So finally just a mention of the conformance change in this release we have moved up to the latest versions of the database platforms that we support both IBM have released new versions of DB 2 and Oracle a new version of that database platform as well. So we are supporting those of this release and probably the most significant changes that we have added some powerful Firefox as a browser as well as Internet explorer. In fact you will notice that Trevor was using Firefox earlier and I can't tell you know that you know there are some performance benefits in the client if I am using Firefox. So if you have the choice in your own environment and that will not be about a way to go. In terms of operating system we support and Windows 2000 and in this release on the service I mean we have extended as a powerful visitor to all clients. So adding to the web clients we have also added support for manager for analyst and for the contributor administrator console and we have for the Excel add in, the export to Excel functionality, we have also moved up to the latest version of Excel. So that concludes this presentation some very exciting features some great performance improvements as well. So got there and I hope you all enjoyed using IBM Cognos plan in 8 version 4. Thank you for your time and enjoy the new release.