

Joseph: Hi. I'm Joseph Pusztai, the Program Director for IBM Cognos Mid Market Products and in this session we're going to get an in-depth view of the in-memory analytics capability of Cognos Express. The analytics server is available to every product within Cognos Express – Reporter, Advisor, and Xcelerator. We will be using the Xcelerator product to, for the most part, to demonstrate the In-Memory Analytics capabilities of Cognos Express. It does give us a lot of flexibility to look at the models that you can build with the In-Memory Analytics Server, as well as the flexibility of the Excel based front end to build business optimization applications. So you can use this In-Memory Analytics Server in a number of different ways. In this demonstration, we're going to be looking at its use for profitability analysis. So we've been looking at a window here called Server Explorer which is our management console if you will, for the In-Memory Analytics Server and the most important objects within this server are something called cubes. As you can see the name of the server is CXMD. That actually stands for Cognos Express Multi-Dimensional. I think one of the most important attributes of this analytic capability is its multi-dimensionality. The vast majority of business tasks that you undertake have many dimensions – customers, products, time, scenario, versions, on and on and on – and when you are working within Excel you tend to get tripped up by the fact that Excel is de facto at two-dimensional workspace. So what one of the key benefits of the In-Memory Analytics Server is the fact that it's in memory, first and foremost. It runs in memory and therefore it's very fast and you can do a lot of high performance what-if scenario type of analysis, as you'll see through the course of this demonstration. The other is that it's in memory and it's also a multi-dimensional in-memory analytics capability and hopefully through the course of this demonstration we've kept it very simple, I'm showing you a very simple model here. Your models, will invariably be more complex and demanding. And of course the more – the harder the model is, the more benefit you're getting from that technology. Simple things, you can certainly do an Excel, but as you move up that complexity curve and move up the data volume curve, you're going to get hurt more and more if you try to do it in Excel. So, In-Memory Analytics Server of Cognos Express lets you do almost anything you can think of in terms of multi-dimensional modeling. Okay, so what we're going to do, we're going to go through these cubes, so as you can see, there's four cubes in this model and I'm going to introduce you to them one by one so it'll make sense as we go through the rest of the demonstration. I'll start off with this cube called Product, and if I expand that, you'll see there's two categories underneath it – Dimensions and Views. So a cube is basically a collection of data that is defined by some number of dimensions. So as you can see here there's in fact exactly two dimensions that make up the product cube, SKUs, standing for Stock Keeping Unit. We're going to be working at that low level of detail, very common in the manufacturing and retail industries. But a good proxy for any kind of detailed data that you might want to manage, be that customer accounts, what have you. As soon as you move beyond sort of the general ledger account cost center level of complexity the ability to do these very detailed granular low level models is quite powerful. So if

I open up, I can actually inspect what that dimension looks like just by double-clicking on it and as you can see here it's a hierarchy of those SKUs, those Stock Keeping Units, organized into product lines. So if I keep drilling down on that, you'll see that this is a very, not arbitrary, but an artificial model here. Line 1 consists of SKU 1 through 100(?), line 2 consists of SKU 100 through 199 and so on and so forth, all the way down to product line 500 which has the final. So, when we add it all up, we've got 50,000 SKUs that we're managing in this relatively simplistic hierarchy here. So that's the SKU dimension. It makes a couple of appearances later on. The things that we're measuring for those SKUs, again, very simple. You can think of this product cube as your product portfolio. I'm simply storing List Price, Standard Cost, and Weight Class. So let's take a look at the cube itself. The cube data, so far all I've shown you is the dimensionality of that cube. Now we're going to actually open up one of these views and take a look at the product portfolio by SKU, so I can simply open that up by double-clicking on it. And this is what a cube actually looks like once it has data in it. SKU 1, there's my list price and my standard cost and you're probably wondering at this point what is this weight class all about. So this model is a Profitability Analysis Model for a manufacturer let's say and the products that they build, every SKU has – is assigned to a weight class. Again, we've kept it very simple. There's only three weight classes – 1, 2, and 3 – and you can see there's a fairly random distribution of those if we scan your eye on it. And the idea behind it very simple, that the heavier an object is, the more it will cost to ship and distribute and as a result of that, it will change my profitability picture by attracting a higher amount of indirect cost. So as you'll see later in the model, I'll have all sorts of ways to play with the weight class and the different scenarios of the weight class, but the first order of business here in putting together this analytics model is to assign a weight class to each and every SKU and they would typically come out of your operational systems, your ERP systems and your manufacturing systems, whatever the recordkeeping systems of your business are, those attributes of your products would be resident there and we simply pull those out and I'll show you that as well how we get access to back office operational data to build up these In-Memory Analytics models within Xcelerator in Cognos Express. So that's the product line – product cube, excuse me. The next cube I'm going to introduce you to is something called the Indirect Cost Factors. Again, we kept it very simple, but if I open up the view onto that cube you can see what I've done here is by weight class 1, 2, and 3 I have three named scenarios – Optimistic, Nominal, and Pessimistic. So the idea here is that my cost accounting team will go in here, look at the distribution costs. Historically, make future projections based on fuel costs and whatever else is driving the distribution cost within my business and basically come up with three different scenarios. My nominal case, this is kind of what I'm – it cost me \$5 to ship a Weight Class 1 and \$10 to ship a Weight Class 2 and \$15 to ship a Weight Class 3, but if I want to do some what-if analysis, if fuel costs go up or fuel costs go down, what would be my expected change in my profitability as a result of that. And as you can see here, the data is highly fictitious but the important thing to remember is that this

model is kept simple for a reason. It's to gain understanding into the capabilities. One of the capabilities is for you to define as many of these cost scenarios as you need to. We've defined three here, but it's not unusual to see in-memory models based on this with dozens, sometimes even hundreds of named scenarios. The reason you can do that with Cognos Express is because it's an in-memory analytics server and that means that as you go through and work with the model you basically pick one of these scenarios to work with and the actual business calculations are only done on demand in real time so it's not as if you would have to go through and calculate every possible permutation and result in a huge amount of data. It's actually got a name. That phenomenon is called Data Explosion, which we avoid very nicely, because it's in memory, we're able to simply do things on demand, on the fly and not basically clutter up your workspace with a lot of extraneous pages that you're not interested in. So that's the Indirect Cost Factor cube. The third cube is a very, very simple cube. It basically tells the model what cost scenario is currently in effect. So right now, if I were to calculate my profitability, I would be using the optimistic cost scenario and you'll see later how that is actually invoked. So those are the sort of three helper cubes. The real heart of the model is this cube. As you might expect, it's a profitability model. It is represented here in this Profitability cube and I've got a couple of different views on that. Now you know how views work. I'll show you sort of the raw data. The only thing I need to load into this cube is my units sold by week and it'll open up that view and you'll see those product lines that we saw earlier – Product Line 1, 2, 3 and so forth. Units Sold is what I've loaded. So every cube in the In-Memory Analytics Server can be a combination of physical data that you load into the cube from your geo let's say or from your ERP systems, and then any number of calculated values based on that raw data that you load into the cube. So this cube has both and the units sold is the only thing that we're actually loading. Everything else in this menu that you're seeing, total revenue, total direct cost, gross margin, etc., etc., is calculated on the fly as we go through and make changes to the data and make changes to the hierarchies and make changes to our assumptions and drivers and factors. So this is physical data and if I switch now to the profitability view on this cube, you'll see in a moment what I'm talking about. So there's my units sold and now my total revenue, total direct cost, gross margin dollars, gross margin percent, indirect cost, net margin, and so forth are all calculated dynamically based on business rules and if in fact I click on one of these cells, I can trace the calculation and see exactly what the calculation engine is doing and you can see there that net margin is a function of net margin dollars divided by total revenue. That's the auditing capability within this server explorer environment. How you actually define that rule – if I just move this over slightly, you'll see this little calculator object, if I open that up, that's basically what we just saw in the calculation tracer. There in fact is the rule statement that is being invoked in that view. And every one of these, total revenue is basically, you know, I introduced those helper cubes initially because they are now playing a part in our model, so the total revenue in this cube is the units sold that we've loaded and then it basically does a reach into that product cube and goes and gets

the list price. And for total direct cost, same idea, units sold, reach into the product cube and go and get the standard cost. And then the gross margin calculations are fairly simplistic here. And the only one that's got a little bit of complexity to it is this indirect cost factor. You'll notice here, I'm getting the cost scenario, the effective cost scenario, using that to look up the actual distribution cost by SKU and calculating my indirect cost and once I have that, basically, I can calculate my gross margins and so forth. So that's the rule statement. And again, this is an In-Memory Analytics technology. It is very simple for you to create new rules or to modify rules if for example I wanted to tamp down some of this indirect cost factor, I can simply come in here and say it's 0.9 percent of that, whatever that factor is, so a lot of power at your fingertips here to change the business rules dynamically. Pretty much the same level of flexibility that you would have with doing this within Excel if you were to use Excel to do this type of modeling, with the exception that all of the heavy lifting is done in a much more elegant way. All the 50,000 SKUs and all their hierarchies and all these calculations are defined in one place and one place only versus Excel as you know when you start trying to manage hierarchical information and calculations across time then you're stuck with this unfortunate copy and paste scenario that is a one-way street, once I do the paste, nothing tells Excel that it needs to go back and update anything versus if I change this business rule here, Cognos Express knows exactly what needs to happen. So it's asking me if -- it thinks I edited that rule and I'll just say, no I don't need to change it. Okay, so now you've seen all the pieces parts of that model, almost, there's another very interesting, important aspect of the In-Memory Analytics Server, as I mentioned, the question mark in your mind might be how do I actually get data and my business information into this model, and you do that through something called Processes and the full name of this subfunction within the analytics server is called Turbo Integrator. As the name might suggest, it helps you integrate data within Cognos Express and it's fast, so once again, the in-memory aspect of the analytic server starts to provide benefit. The load times are very quick. As you'll see, I'm able to rearrange 50,000 SKUs into new hierarchies very quickly and that's exactly what I'm going to do now is show you how simple it is to switch from one hierarchy, let's say my primary hierarchy, to an alternate hierarchy. As you saw here, let me just go back to remind you what my hierarchy looks like right now. So if I collapse that down, you'll remember Line 2 is made up of SKUs 100 through 199, so basically 100 sequential SKUs per product line. So what I'm going to do is go back and update that. And if I open that up, you'll see that I can reach into a database record here. So I've selected, this is a very simple SQL statement that goes into wherever this hierarchical data is being managed. You can manage that directly within Cognos Express if you like. You can basically right click on a dimension at any time and edit the dimension structure if you choose to. In practice, dimensions that have 50,000 elements with multiple hierarchies, it's usually better to do that in an external database of some sort. In fact, there's a whole category of products called master dimension management products, which in fact Cognos has one called Business Viewpoint that is compatible with Cognos Express. But for now, for

demonstration purposes, we have a very simple table that I've created externally that I'm reaching into to pull this information in. so that's the process. To execute it, I simply click on the lightning bolt and you'll see it kind of zips through 50,000 records, rearranges the hierarchy. And to verify that, I can just come back here, recalculate my Profitability cube and that'll take a second or two as it recalculates everything so it's real-time, but it's not instantaneous. And if we expand the Product Line 1, you can see what I've done here. It's kind of an interesting pattern that I've generate, again, using test data, but you can see Product Line 3, 4, 5 – it's not groups of 100 anymore, it's this every other, every 500th element basically SKU that gets rolled up into the different product lines. Obviously, there would never be this dramatic change in your product hierarchy and your business from one cycle to the next, but regardless, if there is and you want to do these kind of massive what-if scenarios, all we did was change the hierarchy, everything else, all the calculations that would say, hey, is Product Line 1 more profitable than Product Line 2. Based on this new mix and this new structure, it's all done dynamically. So that's the back end of the model, so if you were to build this, you would basically be looking at structuring your data into some number of cubes, determining what the business calculations are that you need to focus on. But as an end user, you probably would not want to be working within this environment all the time. This is more of the management and design environment of the In-Memory Analytics Server. The way we actually present this information is, well, certainly in an Xcelerator product, is through Excel as the name implies. So what I've done here, I've built a front end, if you will, a presentation layer, to this model that we've been looking at for the last little while and what this front end does, it lets us basically operate the entire model through a single worksheet, play what-if analysis and right here now, I can start making changes to those assumptions. So let's say my fuel costs have gone up and this is now \$10 and this is \$25, we're going to actually do something dramatic here so you can see – and I'm going to recalculate this workbook. Keep your eye on the top 10 ranking there and you should see some subtle jockeying for position. Okay, you saw a couple of product lines change. Product Line 4 is still in the lead with a 75 percent net margin, but the idea here is that a business analyst can play these what-if scenarios, make wholesale changes to the scenario that's in effect. So you remember, we had three scenarios, so I can very simply just switch this from optimistic to nominal, recalculate and let's see if that has any impact on our relative ranking. You should see some subtle changes there. Yeah, some of the 448(?) is now on the top 10 list. One of the things I've left off on purpose is how you would provide users the ability to invoke those Turbo Integrator processes that I showed you to change the actual hierarchy so I've been modifying the various drivers by weight class. I've been changing the scenario that's in effect. I can switch to an entirely different month. Let's see what my profitability picture looks like here in February and recalculate that. Let's see if line 4 loses its number one spot in February. No, but there's been some other changes. And you saw some of the profitability percentages change as well. How I would do that is well, I pick a spot where I'd like to include those buttons to run and I'm going to put a

little hierarchy to use. Okay, so then for my – from an end user perspective, what I do is something called inserting an action button and I can see here it lets me define which action buttons I can – I’m sorry, which Turbo Integrator process I’d like to invoke and then I can change the way that looks – primary, change the font a little bit, say Okay. And there’s the primary. And then I’ll do the same thing again over here. We’ll insert another one. And this time we’ll take the alternate and say that’s alternate _____ and as I can see, I’ve got to make this appear any way I want and there’s the alternate. So, now my end users have an ability to very simply switch between one and the other. And now I believe, if I run this you will see line 4 finally get knocked out of the number one spot. I will invoke the primary hierarchy. Remember, this is the SKUs 1 through 100, SKUs 101 through 200 and so forth and you’ll see Turbo Integrator now zip through those 50,000 records, rearranging things on the fly and come back to me with a report of its success. And now all we’ve done so far is change the actual hierarchy. What I need to do is recalculate this workbook to see the impact of that and actually, it’s part of the Action button. You see now, line 4 is no longer even on the list because of this fairly massive change to the product hierarchy. Now line 466 is the most profitable line. So in one single workbench we’ve created a very nice, easy to use interface where an analyst can modify all the drivers and factors of their business. They never need to know what’s going on behind the scenes inside the In-Memory Analytics Server, but depending on the skill sets and organizational roles, you can decide to provide sort of a bolted down interface like you’re seeing here strictly for an analyst or some combination where they can go in and look at those cubes directly. So we’re close to the end of the session here. What I’d like to do in closing is give you a feel for what-if I have just some data out there that I’d like to analyze in my – in the In-Memory Analytics, how do I actually go out and get some data and do some quick analysis on it? So the way I’d do that is with a component that’s also included with Cognos Express called Cognos Express Data Advisor which lets me very simply go in to a database of my choosing. In this case, I’m going to use a database that comes with SQL Server called North Winds, it’s a very typical product suppliers customers and so forth data model. So what you see here is the actual database that is made up as I mentioned of employees and order details, orders, and so forth. You’ll see some advisories here. It says add at least one column(?) to dimension values and assign a date/time field. Well, we’ll get to that in just a minute. So what I need to do is basically pick the tables in this model that I’m interested in. In fact when I start off, they’re all selected, so the first thing I would do is start unselecting them and if at any time I’m kind of confused as to what data is in here, like for example regions and there’s also a table called Territories and Regions. I’m not sure which is which. I can simply click on the Preview button and see the Region table there. I can see there’s four regions and if I click on Territories, I can see – oh yes, there’s territory must be a child of region because there’s more territories than regions. If I kind of want to do a bottom up versus a top down, I can disable all the tables and then just start selecting the tables I want. And once I get some proficiency with the database, I can probably work just at the table header level. I

don't have to look at the whole column, I can select and deselect and reselect columns within the table, but in this case, I'm just going to look at the tables themselves. So we're going to do a simple analysis here. We're going to use Products, we're going to use Order Details, we're going to use Orders and we're going to use Suppliers. So those are the four tables that I'm interested in to do some analysis on. My next step here, and there's only three, the first step is selecting the data I want to work with. The next step is defining the relationships those tables have to each other. And if you'll notice here, Data Advisor has already made a guess as to how those tables are related based on the information that is stored in the database. If it doesn't guess right, you can always override it or create relationships as well by dragging and dropping column names to form those relationships. Also if you forgot a table, you could always just go back and say, oh, I really needed the employee table and you should see that kind of just pop right in there. And it finds its relationship to the Orders table here, employees are related to orders through that field and in fact, employees are related to each other. If you scroll down here, you'll see that there's a "reports to" relationship here which it detected. So we're actually not going to use that. We'll stick to our original plan, but it's that simple to go back and pick new tables. So, if the relationships look right and you haven't forgotten anything, you can now go to the third and final step which is now defining what the analysis cube should look like. So the first thing I need to do here is pick something to measure. So add at least one column to dimension values. Okay, let's look at quantity. So I'm just going to drag that in here. And then the final advisory, it says assign a date/time field to dimension time or disable dimension. So I can say, I'm not doing any time based analysis here, I'll just disable it, but if I do enable it, then basically I have to tell Data Advisor what date to use here. I drag it in. And now the Preview button does something slightly different. You'll notice here it's giving me sort of a template for how to analyze data. We talked about the In-Memory Analytics capability of Cognos Express, this is really where it shines as well as the ability to go down really to low levels of granularity, if you need to. If you don't, you can kind of stick to traditional years into quarters into months and in fact, that's the preview of what that dimension will look like when I actually build a cube. And the last thing I need to do now is, I've got the things I'm measuring and the time, the field that I'm using to do the measurement, now what I'm actually going to measure here, I'm going to look at products, and what company supplies those products to me. You'll notice if I drag, depending on where I drag and drop, different things will happen. If I drag here, what Data Advisor is suggesting is that this is a hierarchy of company name and product name, so supplier supplies certain products, let's take a look at that preview. These are all the suppliers. If I expand Big Foot Breweries, sure enough, that looks like products that a brewery would supply of some lager ale stouts(?), _____, these are different Italian cheeses and so forth. Yeah, this looks great. This is exactly the analysis I want. And now I can save this definition off and come back to it at a later time. Perhaps you might give it to a colleague if you're a little bit more technically astute than they are and they can use that and do their own refinements. But, you probably want to just click on

the Analyze button and there is Cognos Express again asking me to authenticate to the Express server, which I will do and at this point the – a number of things are happening. In the earlier part of the demonstration I showed you those Turbo Integrator processes, effectively, this is doing that on your behalf. It's figuring out what SQL queries it needs to run, and Turbo Integrator processes it needs to run to create the dimensions, create the hierarchies, create a cube and then finally when it's finished that, execution of that process. It will present me an Analysis view. Now it's using exactly the same In-Memory Analytics Server as Xcelerator was using but it's going to present this information to me using Cognos Express Advisor which is the analysis and visualization product of Cognos Express. So it comes up with a default analysis and now that you know what cube views are all about, you know that you can simply kind of pivot, drag and drop, expand that down, do some graphical analysis, change the various chart types and so forth. Just to close off the entire demonstration here on In-Memory Analytics, if I go all the way back to server explorer, do you remember we started with these four cubes that were the only cubes in my model. Now if I refresh the display basically, you'll see that here's this new Cognos Express Advisor. My analysis cube sitting right in there and those are the dimensions that we had built as part of the Data Advisor thing and what's really nice now is I can simply use exactly the same Xcelerator techniques that I've used in the past to build additional analyses right here within Excel and maybe extend that into some dynamic planning model and so forth. So all the different modules as kind of the presentation layers available for Express, the Reporter Interface, the Advisor Interface, and Xcelerator Interface are really all just working on that single In-Memory Analytics Server. Thank you very much for your time.

END