



RESOURCES

Home
Resources
Resource Center

OPUS Platform Foundations



Reference

Link: <https://www.tracelink.com/resources/tracelink-university/introduction-opus-foundations-presented-futurelink>

Bob Sturim: I'm really excited today to be talking to you about the OPUS platform. For TraceLink, this is particularly exciting because it's the culmination of over seven years of work to get the platform to the point it's at.

There's a lot to the platform. Today, we're just going to touch upon some of the core foundational elements of it.

Starting with what differentiates the OPUS platform from other platform as service providers in the industry is that OPUS is the only platform that is specifically designed and architected to enable these multi-enterprise applications and solutions, which are critical to the processes that Burke talked about in the previous session.

The ability to digitize your supply chains to be able to collaboratively, along with your trading partners, participate in an orchestrated business process.

Key to that capability is the fact that the OPUS platform is built on top of the TraceLink network, which is represented on the blue box at the bottom of that diagram. TraceLink, through the network operations team, has onboarded approximately 290,000 companies to the TraceLink network.

As Burke mentioned, each one of those companies has been vetted and authenticated so that you know when you link to a partner on the network, you're really linking to that partner, not someone who happened to grab their domain from their basement.

This enables capabilities that are truly distinctive. For example, if you think about how most traditional SaaS-based applications work, they work on a traditional multi-tenant model. Multi-tenancy means that the SaaS provider can enable multiple customers to share the same physical machine hardware, the same physical storage devices.

Multi-tenancy means that each of those customers operate within their own tenancy. They operate as an island unto themselves. Multi-tenancy ensures that what customer one is doing has no visibility to what customer two is doing. That's a great way to build certain type of applications. It's a terrible way to build other types of applications.

Imagine, if you will, for example, if you want to build a social networking application using a traditional multi-tenant model. This would mean that if you're connecting, you would have to log into one account to join your friend Fred's tenancy so you can share cat pictures with

Fred.

You would have to log out of that account, log into a different account to share pictures of your niece's wedding with your friend Jill.

TraceLink is architected in a manner that's much more analogous to social networking. It's built on a network. Because it's built on network, just like that social networking app, you can now connect, you can link to any of those 290,000 companies on the network to begin sharing in these multi-enterprise collaborative orchestrated business processes that cut across corporate boundaries.

We'll talk about some additional foundational concepts, including applications built on top of the OPUS platform, solutions that are built on top of those applications, the integrate capabilities, as well as process networks.

Starting with applications. Applications provide the core functional building blocks that enable customers to perform these orchestrated processes on top of the OPUS platform. The applications are standardized.

What do I mean by that? It means that all TraceLink customers who choose to use that application are using the same code base, using the same version of the application. Customers do not customize applications.

The applications are also headless, which means that they don't ship with a user interface. They ship with a set of APIs. We'll explain why in the upcoming slides.

There are, broadly speaking, two classes of application on the OPUS platform. There's enterprise applications and there's multi-enterprise applications.

Enterprise applications are applications that are primarily intended for you and other employees of the company that you work for. You see a couple different types of enterprise applications listed in the diagram. We'll talk about some of those in the upcoming slides. I'll focus, for example, on the administration app in the bottom right-hand corner.

Administration is the application you use to be able to enable employees of your company to have user accounts on the OPUS platform as associated to your company. It's where you configure the applications that you choose to license. It's where you establish links to your trading partners in the context of those applications.

It's where you assign membership, which users can participate and use those different applications. If you think about it, it wouldn't make any sense for your trading partner to have access to your administration. That's for you to manage your company.

Also, your trading partner has their own instance of the administrative app to administer their company. Those are enterprise applications. In contrast, multi-enterprise applications are applications that are specifically intended to enable you to participate in these multi-enterprise collaborative orchestrated business processes with the trading partners with whom you choose to link.

What you see are two of the first multi-enterprise applications, including APT, which stands for Agile Process Teams. This is an application that allows you, collaboratively with your trading partners, to manage a set of business documents through a defined life cycle of workflow states.

For example, in APT, you could log a quality issue if your supplier shipped you a set of bad parts, and you could assign that quality ticket to your trading partner. Then you could, collaboratively with your trading partner, work through that quality issue through the life cycle until resolution. Both you interacting with the same ticket, looking at the same source of truth.

MPL, which stands for multi-enterprise process links, is the application that provides the foundation for the MINT solution, about which you'll be hearing quite a bit today.

Now, on top of the applications are the solutions. I mentioned before that the applications don't ship with the user interface. That's because the user interfaces are included as part of the solutions. What are solutions?

Solutions are a layer that sits on top of the applications. They're specifically intended to allow customers to extend and tailor the behavior of those applications to suit the needs unique to their business and their set of business processes.

Solutions include a number of different core capabilities. They include, as I mentioned, the user interfaces. Through the OPUS Solution Environment, which you'll hear a lot about today, the no-code designer that allows you to create whole cloth new user experiences, or tailor the default user experiences that ship with the product.

You can define your own set of screens or tailor existing screens. It includes the workflow. Workflow is essential. Workflow is a set of defined states that manage a document through its lifecycle.

OPUS applications ship with a standard set of workflows, but you can extend those workflows as part of the solution. For example, maybe you have a workflow and the document enters a review state. In your business, review means it has to be reviewed by department A, and it has to be reviewed by department B.

You could take the standard workflow that ships as part of the application, and you could extend it. You could say within the review state, there's two sub-states, department A review and department B review.

When a document enters the review, it'll enter the first state. It'll be reviewed by department A. They'll review it, they'll approve it, it'll go to department B. They'll do the same thing. When they're done, now the workflow moves on to the next step.

You can use workflow to extend the behavior and the logic of the applications to suit your specific business needs. It's important because that's the primary mechanism that you have available to you to inject your own logic into the OPUS applications.

OPUS applications sit with a standard set of business objects. The objects, the documents, that you are using to collaborate and manage these orchestrated business processes.

You can take these standard business objects, you can extend them to add your own fields and attributes, and complex data structures, if need be, to suit your own specific business needs. Includes the menus, how users interact with the system, navigate around, which screens the users have access to.

It also includes roles. You can define the roles. How you want to partition up your user base so that you can control which users have the ability to perform which operations, which users have the ability to see which data. All that is designed as part of the solution.

What makes this powerful is that solutions is really a first class concept in the OPUS platform. Contrast that to how many other SaaS providers allow you to extend their behavior. In other applications, you go and you customize the code. There's a couple problems with that.

First of all, when you go to customize the code, it can be very complicated. You need very skilled individuals to be able to do that.

Second of all, if you customize the code and something breaks, and you call up the vendor and say, "Hey, my process isn't working," and they say, "Well, you broke it, you fix it." When you start customizing the software, it makes it very difficult for those vendors to support

you.

Thirdly, once you start customizing the software, it makes it very difficult to take new versions as the vendor puts out a new version. Anybody who's participated in a project where you have a multi-year project upgrade SAP to the new version, you know how difficult that can be.

In OPUS, first of all, we've built the OPUS solution environment that allows you to create these solutions using no code. Second, when you develop solutions, we still support you because there's that strong separation, it's clear what's part of the application. It's clear what's in the solution, and you could always take new versions.

As TraceLink improves and has new capabilities to applications, you can always take advantage of those new applications.

On top of the solutions are process networks. Process networks enable you to categorize your supply chain partners in ways allow them to share worksheet business processes. Why would you create a process network or multiple process networks?

Imagine, for example, that you're a pharmaceutical company, and you have a set of CMO partners, and you have a set of downstream distributors. The processes that you follow may vary, depending on whether you're dealing with your upstream CMOs or your downstream distributors.

You might have different business objects, where you've extended in different ways. You might want to have tailored user experiences, where what information you show and how you refer to it may vary, depending on whether you're talking to your CMOs or you're talking to your downstream suppliers or distributors.

You may want to have extended those orchestrated business processes via workflow to have slightly different logic or workflow steps. You might want to have different access controls and how you divvy up which users are allowed to do different things. It's up to you.

One, you can define multiple process networks. When you link to a partner, you link to a partner in the context of a process network. You can design which partners participate in which process networks, and most importantly, you can define what solution each process network uses.

That means when you navigate and start using one process network, you could have one set of user experiences participating using one solution. You navigate to a different process network, you're using a different solution that could be tailored for that specific set of partners.

Basically, process networks allow you to define a network of networks within the context of your application.

Ensemble refers to the visual paradigm that users use to interact with the TraceLink user experience. It's designed specifically to be highly optimal to allow you to operate across multiple process networks.

There's a couple sections of the page. Across the top, you see the user context. This is where you'd log out, search user preferences, access TraceLink support, access the help system, but I'll spend a little bit more time talking about that dark green bar. That's the navigation bar.

Included in there, you see a drop-down that currently says global supply network. That's where you select which process network you're operating in. You can participate in multiple process networks. Some process networks may be owned by your company. Other process networks, your company may be a partner or someone else's process network.

You can easily switch between those different process networks simply by selecting the appropriate value from the drop-down without having to log out, log back in, have different URLs, different passwords, any of that.

Next to that, you see another drop-down. That's the partner drop-down. If you are an owner of a network, you could have a wide lens.

You could be saying, "Show me all the information across all my partners in this process network," or you could choose to shine a spotlight on a particular partner by selecting that partner from that right drop-down. Now, it will only show information related to that single partner.

Then on top of that, you see a tab, user interface, which is also interesting. We support multiple different tabs. Each tab could be operating in the context of different process network, which means that you could easily switch from one process network to the other, and switch back again simply by selecting different tabs.

When you switch back again, OPUS is going to remember what you're doing and reconstitute that page where you left off. Moreover, OPUS remembers what you're doing.

When you log out at the end of the day and come back the next morning, it'll reinitialize the tabs you had open when you left the night before and remember what you were doing in each one of those tabs when you logged off the previous night.

Along the left-hand side, you see the navigation pane, the left-hand navigation pane. This is where you select which orchestrated process you want to be operating in. Generally speaking, the orchestrated process are identified by the business object upon which that orchestrated process operates.

You'll notice that most times what you source up on the left-hand side are a set of different business objects. You choose the business object, and that's the process that you're zeroing in on.

On the right-hand side, you see what we call the push panel. It's called the push panel because you can hide it. If you choose to show it, it'll push out from the right-hand side. It doesn't hide anything. It doesn't block you from doing anything. It's non-modal.

You can easily keep doing whatever you're doing. Depending on what you're clicking on, what operation you're performing, the content of what shows up in the right-hand panel will vary.

For example, what you see here is a search page. If I have the panel open, if I collect different rows, the right-hand side will likely show me additional information about the row that I selected. I can easily select the line and the right-hand panel will adjust accordingly.

In this case, we're on a search page, and you see the filter panel. This means that the user selected the filter button in the top right-hand corner, which has brought up a set of search criteria, which would allow the user to control which set of records show up in that search table.

Integration is another one of those foundational capabilities you had heard Burke talk about in the previous session, which I think is truly unique to TraceLink. As Burke said, we've invested in an Integrate Once architecture.

This means that you can invest to integrate your business processes into OPUS. Now, by virtue of doing so, you can participate in that orchestrated process with virtually any other company on the TraceLink network simply by linking with them without having to change your integration at all.

There's a couple of key components to this. One is that you can speak whatever protocol you

want. You want to exchange messages with AS2, you want to use FTP, you want to send email using SMTP, you can use any one of those protocols, based upon what makes sense for your business.

If you don't want to do any of those and you still want to be able to integrate, you can use link actions. You'll see one of the enterprise apps is the XT link actions app. This allows TraceLink to connect to your system via the published APIs of your system to be able to integrate without having to install any software or customize your system at all.

Or if you don't want to do any of those things, you can just integrate using the UI. You can log into the UI, for example. You can see the purchase orders that someone has sent you. Straight from the UI, you can choose which purchase orders you want to acknowledge.

Additionally, as Burke had hinted at, you can use the message format that you want. You're an EDI shop, maybe speak X12. Some of your partners speak EDIFACT. Other partners are SAP IDoc shops. You could have others that want to send EPCIS message or a custom format. It doesn't matter. This is achieved through what we call transforms.

Another one of those boxes you see is a transform service. A transform in OPUS speak is a small program that converts from one data format to another. When you integrate to TraceLink, you configure which set of transforms you want to use for which transaction.

What this means is, again, if you're that EDI shop speaking X12, you would select to use transform specific to X12 messages. You send a PO, for example, into MINT, and it would go from X12. The transform would convert it from that into the canonical format, which is the format that TraceLink applications think natively.

Then TraceLink may forward that PO to your downstream supplier partner, and that partner may have configured that they want to use the SAP IDoc transform. Again, SAP TraceLink will convert that format into the IDoc compatible format, and send it via the protocol that partner has selected.

The main thing with all of this is your integration does not change at all, regardless of the mechanism that your partner on the other side of that exchange is using. When you're operating within the MINT solution, you use the same orchestrated business process without regard to the mechanism that you or your trading partners use to integrate.

It allows a seamless and holistic view of your data, allowing everybody to connect using the technology and formats of their choice.

Lastly, you see the integration catalog. This is a searchable catalog. We'll talk more about that in a slide coming up. It allows you to peruse and find transforms at TraceLink, or maybe other vendors has published.

It allows you to peruse and find link actions that TraceLink, or perhaps other solution providers have published, that allow you to connect and configure which transforms you want to use, or if you're using link actions to connect, you can use those link actions to connect your system.

You can also use those as starting points. If you want to use, for example, a link action that TraceLink has created, but maybe you've customized your system, you need to tweak it some, you can take the link action from the catalog we provide, you can extend it and make it your own link action.

Metadata. Sounds really complicated. Metadata is data about data. For us, it is basically information about the business objects, the application the solution uses.

We have another one of those enterprise application, is an application called Metadata Manager or MDM. It understands all the business objects in the system, not just what the

structure of those objects are, what fields are on the objects, but has a richer set of data.

For example, we have a concept of rich field types. What that means is that you might have a field on, say, a purchase order or some other business object that's a numeric field.

We know it's not just a number that's stored in the database, but what that number represents. Is it a percent? Is it a currency? Is it a date, or a time stamp, or a duration? Is it a telephone number? You define the rich field type.

Why this is important is TraceLink knows how to render the information on the screen, depending on the type of rich field type you've associated with that field. If it's a field that you can edit, TraceLink knows how to present the information, how to validate the information, and enter it.

For example, if you indicate that that numeric value is actually a date, then TraceLink knows to render a date picker in the UI. It also supports operations.

Operations are the actions that you can perform on a business object. The standard CRUD-based operations, but there could be other types of operations, like release something, or publish something, or follow something. All those are defined on the objects in the system.

This allows you, when you build your UI, to select not just what fields you want to show up in the UI, but what operations you want to allow the user to perform against that object in that screen. You can simply put a check marks and choose which operations you want to enable without having to wire those in.

Last, we already talked about workflow. Workflow is stored as metadata on the objects. Workflow, again, is critical, because it's the primary mechanism you have to be able to take the base logic that ships standards part of the OPUS applications, and extend it in ways that are specific and meaningful to your specific set of orchestrated business processes.

All of this allows you to build those no-code user interfaces. We have different sorts of page types for creating new records via the new page, the search page, the view edit page. It's also where you define your reports and dashboards, which you'll also hear about more this afternoon.

In all these things, metadata is essential because it basically allows you to find what you want the system to do, and the system could figure out how to make it work without you having to write code.

On top of the metadata is the OPUS Solution Environment. This is a truly no-code solution builder experience that allows you to take standard solutions that TraceLink provide, or create fresh ones from scratch, or modify the standard ones that TraceLink provided, all without having to write code.

It allows you to modify some of the objects we talked about earlier, those business objects through the OPUS Solution Environment. You can take the standard objects that ship standard from TraceLink, and you could extend them, add your own fields or complex data structures to those objects to make them yours, make them appropriate for your set of orchestrated business processes.

You could define the workflow. You can take the standard workflow that ships with TraceLink, and you can extend it. You can define what sub-states you want, what transitions are valid between those sub-states, and you could associate logic, what should happen under the covers when that document moves from one state to another.

It's where you define the menus. Again, how the user navigates to the system, has access to screens, what screens they have access to. It's where you define the UI, the pages, and it's where you define your roles.

How are you going to control which class of users have access to perform which operations, see which data? All of which enabled through the no-code solution designer, and you'll hear quite a bit more about that for those of you in the solution designer track this afternoon.

Lastly, we have the catalogs. I hinted at catalogs earlier when we were talking about the integrations. Catalogs are standard platform capability that provide a searchable repository for certain types of objects and assets. For example, we've talked about solutions.

TraceLink application ships standard with a standard solution. That solution is fully functional and usable out of the box. If you don't have a need to customize and tailor your experience, you can use a standard solution without modification.

If you did want to take the standard solution and use it as a starting point for your own extensions, you could peruse a standard catalog, find the standard solutions you want.

You could basically do a save as, which would copy that solution and make it an available company solution, at which point it's available in the solution environment for you to extend and tailor in ways that make sense to you.

Then if you wanted to, you could publish that solution to the company catalog. This is a version of the catalog, but it's private to your company. It's not accessible to other customers. If you have tailored things unique to your business, you can put them in your company catalog.

One reason you might want to do that is if you're developing your solution, you're probably doing that in the validation environment. If you publish it to the company catalog, now your production company, your company in the production environment, has access to that same catalog entry, so that's how you'd make the solution live in prod.

Then the marketplace catalog is where TraceLink, or potentially, solution providers will create other assets. We could have third parties, for example, creating solutions specifically targeted to specific segments of the supply chain.

In addition to solutions, this is where you can find link actions and transforms. You can peruse the marketplace catalog for link actions and transforms that TraceLink or other solution providers have created.

You could use those out of the box, or you could choose to copy them down and load them into your company catalog, if you would need to extend them in ways specific to your business.

[background music]

Bob: It's also where you have dashboards and reports, so you'll see catalogs used throughout the system.