One Network Enterprises™

Starting from Scratch: Rethinking the Supply Chain

One Network has delivered and continues to enhance and develop an approach to supply chain planning that we believe corresponds closely to what leading analysts, thought-leaders, and technology visionaries predict will become the new standard.

In particular we agree with the observation that hierarchical and serial planning (where broad aggregate planning is done up-front) is at best insufficient and at worst misleading.

The problems with the traditional hierarchical planning approaches are several:

- **Planning is done at a broad aggregate level.** However, 'reality' is much more detailed and real problems occur at the detailed level. (e.g. Product Family planning vs. Single-item-at-a-single-location planning)
- **Planning is done over a long time horizon.** However, reality is constantly changing and many if not most of the assumptions made are invalidated by the time 'now' shows up.

The reason why planners spend most of their time in spreadsheets doing low value-added work is because their planning tools don't help them with the here-and-now. As a result they are forced to make-do which typically involves a lot of spreadsheets, emails, and calling back-and-forth.

What is needed is a framework that supports near-term-execution as well long-range-planning. However, near-term-execution requires an entirely different conceptual as well as architectural approach.

From a conceptual standpoint, near-term-execution needs to take into account the following realities:

- **The world is constantly changing** and any approach must be able to account for this and not work on an idealized frozen snapshot.
- **Global optimization is neither a desirable nor feasible approach.** Instead the appropriate way to think about this is closer to agent-based planning where agents are responsible for reacting to and optimizing subsets of the overall problem. In One Network terms, we refer to these

subsets as subnets and agents as Intelligent Execution Modules (IXMs).

- **Agents must be able to communicate with each other.** In the One Network approach this communication is done via well-defined state transitions on transactions.
- **Agents must work directly on transactions** and be able to both read and modify transactions just like a human planner would. In particular they should not be working on a simplified snapshot of the world.
- **The micro-optimizations done by agents must be completed in seconds.** If the world has changed (and invalidated some assumptions) since the start of the micro-optimization, the agent must be able to roll-back.
- Since there are too many subnets and agents to manually administer, **it is important that the agents learn from past behavior and modify their behavior (or settings).**

This conceptual approach also leads to a very different architectural approach. Instead of a few sequential, planning runs running on stale data, there are now potentially millions of agents running on the live transactional system. Each of them needs to be capable of rolling back their changes if appropriate. Additionally each must be able to move transactions through their life-cycles. In other words, these agents need to be transactionally aware.

The above set of requirements leads to the need for a horizontally scaling transactional grid type architecture on which millions of agents can be running and communicating. In other words, the architecture must be massively scalable.

While this is a radically different approach to planning, it leads to huge benefits:

- **Planners can focus on true exceptions** rather than running the day-to-day execution. They can also be freed up for more value-added strategic thinking.
- **Micro-optimizations can be done at scale** (for example, for every single item at every location). This is just not possible for human planners to do.
- Because this approach is based on the true-state of the world, there is **no disconnect between planning and execution**.
- **Informational lead times (which are a large part of total lead times) can be drastically reduced** as there is no need

for batch planning engines to 'wait' for the upstream ones to finish.

One Network provides both the transactional platform and the grid architecture to make this planning approach feasible. In addition it also provides over a dozen agent types that do various types of micro-optimizations.

However, One Network has also realized that there is a large diversity of problems and a one-size-fits-all approach is not always feasible. Reasons for this are different transactional processes, different operational goals, different business models, different data inputs, etc.

To this end, One Network realized that it is critical to create a platform on which others can create transactional processes and related agent-based optimizations that are custom to the company in question. This has the added benefit of allowing the company to have complete visibility into the approach being taken (i.e. it is not a vendor-specific black-box).

The One Network platform is unique in that there are a large number of "primitives" that can be reused in creating customized agents. In typical supply chains we have found that 80% of the underlying transactions, processes and agent-logic can be re-used.

With the One Network platform the data model can be extended via a 'transactional polymorphic mixins'. While complicated-sounding, this basically allows multiple applications to independently extend core models so that there is no need to constantly synchronize data, which tends to introduce its own headaches, latencies, inconsistencies, etc. This turns out to be a key requirement for enabling agent-based micro-optimizations.

Extending the data model with various kinds of additional data needed by either planners or by agents without introducing the need for replication or synchronization is key to making the whole approach work.

You can learn more about our platform which we refer to as Developer Network at:

https://devnet.dn.onenetwork.com/oms/apps/DeveloperNetwork/www/