

Propelling digital transformation in manufacturing operations with Power BI

If manufacturers don't ship on time because of a problem in the supply chain or production line, customers and future sales are affected. To stay competitive in the marketplace, prevent costly delays, and avoid errors, it's critical to understand and respond quickly to end-to-end operations in the factory, supply chain, and production lines.

At Microsoft, hardware for products like Microsoft Surface, HoloLens, and Xbox is big business. In fact, in fiscal year 2016, it was a billion-dollar business. This large dollar amount underscores why it's so important to monitor, share information about, and optimize manufacturing operations.

Previously, in one manufacturing facility of Surface, HoloLens, and Xbox, teams reviewed operational information in daily or weekly business reviews. In preparation, executives typically spent hours curating reports from Microsoft Excel worksheets and presenting the data in Microsoft PowerPoint. This effort took valuable time away from more strategic work, like data-based decision-making. Plus, the data often wasn't current and gave different messages.

So Microsoft IT and the Microsoft supply chain created the first phase of an automated solution that uses Microsoft [Power BI](#), [Power BI Mobile](#), and data analytics to help Microsoft and factory managers. Using the solution gives us:

- **Increased productivity.** It's faster to get data and create reports on production, shipment yield, and quality.
- **Improved operation reviews.** We make decisions based on current, trusted, interactive data.
- **Decreased time to root cause, and drive continuous improvement.** We quickly detect process-related issues and anomalies to change how the supply chain's managed.

The results? Even in our initial phase, we're seeing tons of value for executives, factory managers, and stakeholders:

- Time-savings by not creating reports (15 to 20 hours per week, in one case); faster decision-making.
- Significant cost-savings—think of the cost savings of solving a yield problem versus preparing a report.
- Onscreen dashboards in meeting rooms, with data that shows—at any given moment—what to produce or fix.
- Ability to monitor, detect, show, and solve production-related problems hourly rather than daily, from anywhere.

And we'll see even bigger value in the next phase. We'll include more manufacturing locations, collect data from other tier-level suppliers, use predictive analytics, and get more process data from factories. Plus, although we currently support certain Internet of Things (IoT) scenarios, we'll take these to the next level to empower operations even more.

Moving from curating data to responding to data

In [How digital transformation is changing the face of manufacturing](#), Microsoft CEO Satya Nadella states: *"A digital transformation is remaking companies and their factories, bringing the intersection of manufacturing and technology even closer... Enabling that transformation are systems of intelligence that help companies gain insight and take action...optimize their operations, and change the very nature of the business models around their industrial products."*

Supporting the overarching mission of digital transformation, our organization's leadership team wanted to revamp how manufacturing-operations data is used. The aspiration? To respond to data rather than just amass and curate it.

So in Microsoft IT, we started by visiting a contract manufacturing facility (which we'll refer to as CMF1) in China that builds the Surface, HoloLens, and Xbox. We were eager to deploy Power BI, so that factory teams get fast insights on production, shipment yield, and quality control. Having answers to questions like these adds extra value to decisions:

- Are we going to ship late?
- Do we have enough of the necessary products/components to ship on time? Do we have too much inventory?
- Does the quality level meet our targets?

Connecting data and operations

Besides using Power BI to respond to data in manufacturing, we saw potential for using it to form connections—in the data itself, across suppliers, and across factories. For example, we could transform our operations by:

- Connecting the data; connecting data sources.
- Making better connections in the data itself about what the data reveals.
- Eventually using data analytics to see even deeper connections, and driving more predictive and IoT capabilities.
- Eventually linking other factory locations to CMF1, connecting tier-level suppliers, and linking the supply chain.

Deploying our solution, and adding predictive and IoT capabilities

We divided our work into phases. Although we're seeing incredible value and quick wins from the first phase that we've completed, the second phase will build on this foundation.

Phase 1: Connecting the operations data, data sources, and identifying connections

This phase is about connecting the data from different systems and making it easy for factory managers, executives, and others to spot trends and issues.

We set up a model for an initial dataset, and spent a month or so refining the model and the reporting, to align with business goals of factory managers and stakeholders. This data came from Excel worksheets and from our internal tools for managing out-of-box audits (where we randomly select boxed units from the end of the line for analysis and quality check), yield (the percentage of products that are manufactured correctly the first time without rework), quality, inventory, repairs, shipments, and packout (number of units that are built). We want to automate for other business units and include more products and factories.

We also set up large Power BI screens in some of the conference rooms. Today, team members simply walk into the room, and see the reports displayed. And they can also see reports anytime from their office or phone.

Phase 2: Providing deeper data analysis, more IoT capabilities, and more connections

We'll move beyond basic reporting to getting richer insights and value from the data. We aim to be more predictive and detect potential issues before they arise. With machine learning, we'll identify any change or anomaly in the dataset that's further down in our supply chain and likely to cause errors or customer impact. And we'll respond earlier—three to four stages earlier in our process, which helps us save costs and avoid issues.

From an IoT perspective, we see many opportunities and are identifying potential scenarios—such as with our factory partners—considering the machines and operations that we can collect data from. We'll explore using [Azure IoT Suite](#) for these scenarios. Its easy-to-deploy, preconfigured solutions closely align to our vision of:

- Improving operational efficiency.
- Performing remote monitoring and predictive maintenance.
- Ingesting events from sensors in real time.
- Providing a data analysis and visualization framework, and incorporating real-time, streaming analytics.

With IoT, we'll be more proactive and predictive in improving operational processes, which helps us grow our revenue and profits. Right now, our connected devices pull manufacturing test data from factory test equipment. But in the future, we'll bring in process monitoring equipment, and add more data at the process level. Why is this important? Because even before test data reveals any shift in operations, we'll see process steps and new connections earlier.

Last but not least, linking our supply chain is also part of phase 2. We want to connect sub-manufacturing locations with our primary location, CMF1. This way, sub-manufacturers see incoming material and quality. And, if we detect issues in the field, individual component manufacturers, in turn, see if their component is problematic.

Benefits we've gained

Again, much of the transformative value of this solution is still to come with the next phase. But just in the brief time that teams have used Power BI, detecting and solving problems is already quicker, along with other key benefits like:

- **Easy, initial setup.** We set up screens displaying Power BI in CMF1 conference rooms in around six weeks.
- **Significant time-savings.** For the first dataset we worked on, there was a savings of 15 to 20 hours a week from teams not having to curate data. Instead, they just looked at Power BI on a large screen in the conference room.
- **Cost savings.** Imagine the opportunity cost of spending an hour preparing a report versus fixing a supply chain issue that could be worth millions of dollars. Saving 15 hours a week at a high hourly pay rate adds up quickly.
- **More holistic and unbiased data view.** We collect data from other suppliers, like tier 1.5 and tier 2, and have a trusted data view, regardless of manufacturer or supplier—no more datasets with disparate messages.
- **More productivity; fewer meetings.** Executives look at the reports at any time. If the reports indicate that everything is functioning as expected, they often find that they don't need to meet.
- **Cloud first, mobile first—works on Windows Phone, iPhone, iPad, and Android.** Team members get current data on their phones, to make decisions that affect yield, performance, and output.
- **Quick, easy insights with data that tells a story.** We see interactions, such as between yields and returns. No programming knowledge is needed to see inventory levels, process efficiency, costs, machine usage, factory production flow, and material. We detect and intercept anomalies quickly, which prevents issues later.
- **Ability to customize.** We create views that are most relevant for different teams. Also, Power BI is developer-friendly and extensible. Third parties can contribute and create features, graphs, and charts.
- **Integrated data.** In the past, we had to extract data from our internal tools and then integrate the data to see factory production flow. Now, Power BI integrates it for us.
- **Data analytics in Power BI itself.** These analytics reveal areas for machine learning for future scenarios.

Transforming insights into results with Power BI

It's exciting to see the insights that have led to improvements, since the time we deployed Power BI. For example, we:

- Improved yield end to end, and identified the root cause for a particular yield gap. When you produce many units, improving this gap translates into potentially very large cost savings.
- Streamlined inventory and hardware management to prevent aging inventory and ensure the right quantity of products and materials. The result? We've improved excess, obsolescence, and productivity; reduced costs; and improved margins. Again, because of the numerous units, each issue we prevent means millions of dollars saved.
- Saw excess inventory the first week Power BI was deployed, and brought levels back to production targets.
- Saw which line an excursion came from, when, and which material. Factory managers contained the issue.

Figure 1 shows a sample Power BI dashboard—in it, you can see how our factory operations are performing for a specific product. For example, it shows shipment yield, outgoing quality, variance from targets, packout and shipment actuals versus target, repair inventory trends, and other information.

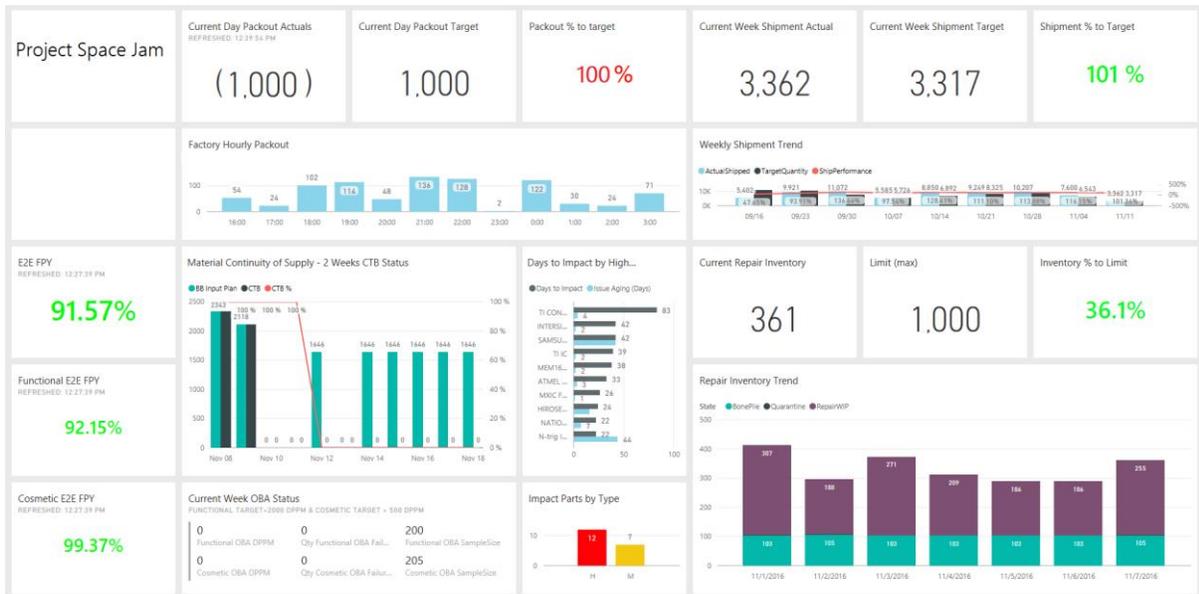


Figure 1. Sample Power BI top-level dashboard—packout, shipment, repair inventory for a particular product

Figure 2 has details on build yield, first pass yield (FPY), weekly/daily trends, and variance from target yields. FPY is the number of units coming out of a process divided by the number of units going into that process.



Figure 2. Sample Power BI dashboard—build yield, weekly/daily trends, variance from target yields

Setting up the data flow, architecture, and workflow

Currently, we're not using IoT Hub in our architecture, but are exploring it for the future. We use service buses, event hub, and Azure Data Lake, but the event hub is unidirectional. Benefits that we could get from using IoT Hub are that it's a bidirectional cloud connection, and it's designed for high-speed, real-time ingestion of events from sensors.

Our test systems talk to each other and to a central hub. We're trying to see what process data we can pull from factory machines, load into our datasets, and bring into Power BI.

Capturing the data

We capture about 125,000 messages per hour, although this rate can vary from about 10 percent to 20 percent. How do we get a single, trusted version of data from the manufacturing shop floor to our database?

1. We have a technical control system in Microsoft—a collection of hardware, software, and sensors—that we install in our manufacturing facilities. Plus, factories maintain their own shop floor control system and data.
2. Our technical control system collects manufacturing data from factory test equipment. It pushes data to the shop floor control system and to a manufacturing test system (which records test yield data from the production line) where we extract the data. There's a local instance of the manufacturing test system at each factory and a centralized instance in the United States, where we collate data from all sources and factories.
3. Data goes to an on-premises source and syncs with a server in Redmond. We want to move the data to Azure.
4. If we can't get data from our technical control system, we check with the manufacturer to see what they collect outside our test system on their shop floor systems. The manufacturer sets up a direct query on their shop floor to go into our data lakes. We want to automate this more, so that data goes directly into our technical control system, bypassing the shop floor.

Designing the architecture and workflow

The technology and process that we used to create a unified system of record for CMF1 includes the following steps:

1. **Microsoft SQL Server 2014, Azure Blob storage, Windows PowerShell.** We moved data from factories in different regions into our SQL Server 2014 central data warehouse and operational data source (for more detailed data, such as the serial numbers of the devices that we've shipped). We used Azure Blob Storage to transfer files from the shop floor control system to the data warehouse, and PowerShell scripts to copy flat files (like those related to shipments, the repair status of defective units, and factory output) from an on-premises FTP folder to Blob Storage.
2. **SQL Server Integration Services (SSIS).** An on-premises SSIS job looks for newly uploaded files and downloads them to the data warehouse. We report directly from our SQL Server environment. We also get data from manufacturing test history. This data goes into the central data warehouse and operational data source.
3. **Power BI and Power BI Desktop.** We use Power BI to display the data and Power BI Desktop to create reports. Our manufacturing dashboard is where we display all these reports.
4. **Master Data Services (MDS).** We use MDS to maintain master data. For example, what operations and lines should we show? When does a shift start/end? If we change metadata values, they'll be reflected in reports.

Figure 3 shows our current high-level flow of factory data.

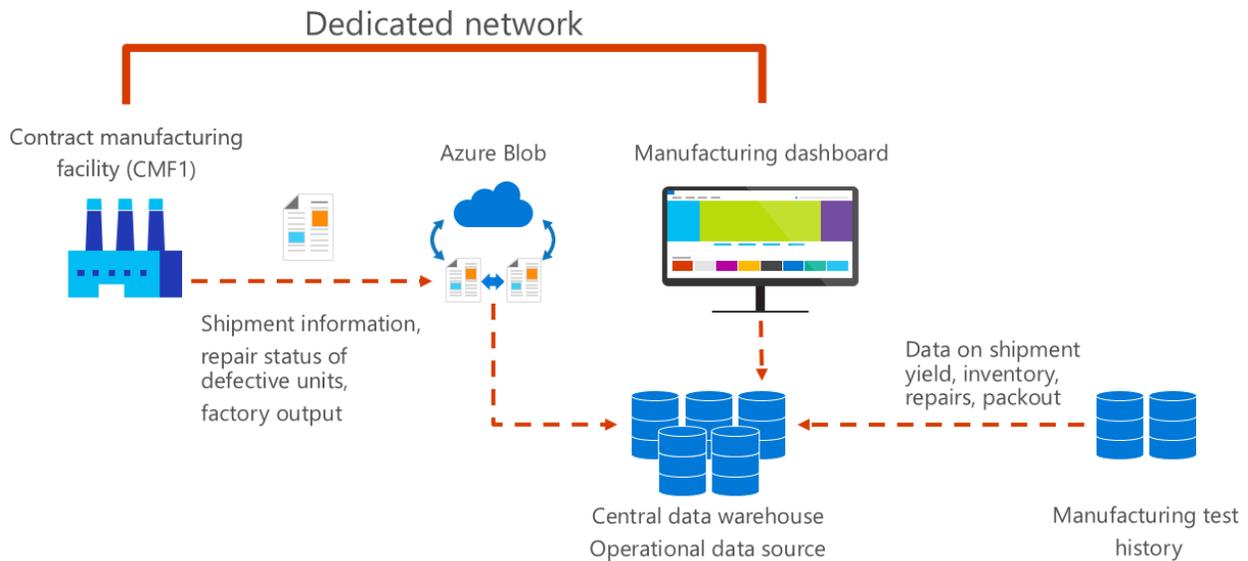


Figure 3. High-level data flow

Providing data governance and data quality

We're deciding on the best way to govern master data. Power BI gives some visibility into data quality, and we also set up alerts in Microsoft System Center Operations Manager to tell us whether we have data in the system. We'd like to get alerts on data quality or missing elements (like part numbers) that would lead to data not appearing in reports.

We're also researching whether we can use statistical process control charts to see if our data is within the statistical quality limits of what's expected, and if we're getting the right data. These charts use statistics and algorithms—based on historic data—to plot data points, trends, and upper and lower control limits. A center line represents the average between the limits. For anything beyond the control limits, we can investigate root cause and make improvements.

Setting up alerting

We're still optimizing alerting. Right now in Power BI, alerts are tied to personal subscriptions. Power BI subscribers get alerts for their own personal rules, settings, and key performance indicators. But they can't get alerts that apply to a user group. We created an alert system outside Power BI, and identified alerts to add at a line-operation level. We're working on pulling that data into Power BI, so that it sends an email alert that points to a Power BI visualization.

Scheduling data refreshes

We don't have streaming analytics. Screens update on their own every three hours, eight times a day. We use the Power BI scheduler to manage the frequency, and eight times is the maximum frequency allowed. On the back end, there are hourly jobs at the factory that push data onto our system, with other data feeds coming in once daily or weekly. Beyond scheduled data refreshes, anyone can manually refresh Power BI data at any time to see the latest dataset. We use Operations Manager to monitor whether we get feeds on time. If we don't receive them, it alerts us, and we then follow up with CMF1 to see why we're not getting the feeds.

Using predictive capabilities and machine learning

Some predictive analytics and machine learning are happening now, but the next phase is when we'll really tackle this and get results. Today, we're collecting data, training the models, and iterating to see what insights are possible.

We plan to move our manufacturing data to Azure Data Lake storage and use Azure HDInsight for predictive analytics and machine learning. We'll see earlier on what's in the pipeline in terms of quality, material, and components. And we'll be able to predict errors and detect critical events that are happening in our technical control system.

One example involves looking at what's happening at tier 1.5 and tier 2 suppliers, and tying yield activities to what's happening at our tier 1. For example, we look for connections between low yield in one part of the process and how it correlates to low yield in another part of the process. After we spot trends, we can set up early predictive triggers.

Our machine learning team is analyzing field data to identify returns and quality issues in the field. To refine our predictions, they're correlating issues in the field with operations in the factory to see the correlations that affect customers. These insights will help us predict and prioritize issues and alerts, based on severity of customer impact.

Technical challenges in implementing our solution

Technical-related challenges in implementing this solution included:

- Finding resources and support for removing hurdles and taking this solution to the next level.
- Versioning in Power BI. When we started deploying our solution, team members sometimes worked in different versions of Power BI, and as a result, they couldn't open files that were a different version. This has been resolved.
- Master data management. We're still working to ensure that the master data is updated and data isn't missing.
- Intermittent issues with data refreshes, latency, and slow performance. Our datacenters are in the United States, and the CMF1 factory managers are in China. Sometimes there are resulting latency issues, which we're working on. We'd like to get access to backend logs from Power BI, proactively monitor when the Power BI service is slow, and alert users. Even without latency issues, some days there's a 50 to 60 percent dip in service.
- We have to manually check for issues each day. Automating this and putting performance information into one Operations Manager alert would help the support team see if the service is slow and take any necessary action.

Lessons learned/best practices

Lessons we learned and best practices include:

- Consolidate spreadsheet data into one location.
- Where possible, automate the data.
- Iterate to identify the ultimate data source—not spreadsheets or email—and link that source to Power BI.
- Fix the data source. Our leadership wanted raw data as much as possible rather than data that's so processed that we can't trust it. Instead of fixing the data, we fix the source so that we get correct raw data.

Looking ahead

We want to build on the quick wins we've seen and tap all the potential value of this solution. We're looking forward to taking things to the next level, so that we can unearth more layers of digital transformation. We plan to:

- Deploy Power BI at other factories globally, and connect other contract manufacturing facilities to CMF1.
- Capture data from other suppliers—tier 1.5 and tier 2.
- Build on IoT capabilities—possibly using Azure IoT Suite/IoT Hub—and add more factory data at a process level.
- Look beyond factories. Integrate information about the status of materials in Power BI, regardless of supplier.
- Predict factory yield and identify potential issues, before they become issues in tier 1 factories.
- Use statistical process control charts to detect anomalies for early intervention, before customers are affected.
- Use machine learning to identify shifts in the dataset that are further down in the supply chain that will likely cause customer impact or error. Adjust and respond to this change earlier in the process.

- Link our supply chain and manufacturing locations. Connect sub-manufacturers with CMF1, so they can see what's coming in terms of incoming quality and material.
- Use streaming data analytics, and capture *real-time* data in Power BI.

Continuously transforming and connecting

By modifying operations in the CMF1 manufacturing facility, we're moving from laborious report creation to quick data in Power BI dashboards. We're connecting operations, preventing delays in the supply chain and production line, decreasing time to root cause, and adding value to decision making. Power BI is disrupting "business as usual" in a competitive market by bringing manufacturing and technology closer together, with data as our compass.

For more information

Microsoft IT

microsoft.com/ITShowcase

[IT Showcase data analytics content](#)

[Microsoft Power BI](#)

[Azure IoT Suite](#)

© 2017 Microsoft Corporation. All rights reserved. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. The names of actual companies and products mentioned herein may be the trademarks of their respective owners. This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY.