FACTORY EDGE-TO-CLOUD ANALYTICS MADE EASY

Simplify Data Access at the Edge enabling IOT-based Use Cases

45%

Industrial products lead all industries in IoT adoption at 45%

The Connected Manufacturing Imperative

Digital transformation is now so embedded in manufacturing, it's hard to find one business process it has not touched. McKinsey reports that $11 trillion in economic value will be delivered by 2025 through industrial digital transformation efforts. The rapid adoption of connected manufacturing initiatives is powered by the following key drivers :

- **Increased Quality** due to both product analytics and real-time data collected and shared in a continuous loop for ongoing process improvement.
- **Improved Operations** including asset condition monitoring, OEE improvement, and predictive maintenance for improved processes and reduced downtime.
- **Increased Agility** by enabling manufacturers to quickly pivot production due to changing market demands or supply chain disruptions.

Manufacturing companies are leveraging data derived from hundreds or thousands of connected assets to drive never-before-seen efficiencies and insights in both operations and IT, leading to entirely new value streams that can transform their relationship with their markets.

Key Manufacturing Drivers

Manufacturers face a myriad of challenges, most of which have traditionally centered around improving quality and reducing unplanned downtime. Recent developments have added agility to the mix as manufacturers strive to rapidly change to meet market demands or supply chain disruptions. The following are driving manufacturers to adopt IIoT solutions immediately:

Quality

Quality control can be traced clear back to the Industrial Revolution, and since then, manufacturers have made consistent improvements in product quality and reliability. Although manufacturers have made impressive strides in quality, competition for market share and the cost of quality, up to a staggering 5 percent to 40 percent of sales, continues to drive improvement. Subpar manufacturing quality also imposes significant downstream costs to almost every aspect of the organization including underutilization of assets, added scrap and rework expenses, warranty costs and lost sales.

Plant Uptime

Unplanned downtime costs manufacturers heavily in both the idle time of workers and machines, thus increasing the cost per finished good. According to Deloitte, unplanned downtime costs manufacturers approximately $50 billion per year and poor maintenance strategies can reduce plant capacity by 5 to 20 percent. Not surprisingly, equipment and plant uptime are a high priority for C-level executives. An IndustryWeek survey of manufacturing executives found the top driver for technology investment is achieving higher productivity through reduced downtime.

Plant Agility

Most companies can derive value from an IIoT initiative immediately in terms of quality and uptime; however, agility has separated the leaders from the pack. The pandemic has dramatically changed almost every sector of the manufacturing industry. Just last fiscal year companies were focused on capital spend expansion and stock buy-backs, while many are now in survival mode and focusing on cash and liquidity management, make or buy decisions, and improving the visibility of ROI in all aspects of the business.

Due to changing market conditions, some plants were able to quickly change from auto component production to medical device equipment, for instance, while others enabled remote equipment monitoring and process improvement to continue plant operations from centralized control rooms, ensuring social distancing. Still other manufacturers benefited from their digital transformation efforts by being able to agilly scale up or down production capacity to meet the changing demand and supply challenges of the market. Their integrated enterprise-wide supply chain and manufacturing operations insight allowed these companies to pivot because of real-time data IT and OT data driving their business.

McKinsey reports that digital transformation efforts focused on agility can have significant returns. Inventory reduction can return between 12 to 30 percent; lead time reduction can return between 10 to 90 percent; and change-over shortening can return 30 to 50 percent. What is the common thread tying these quality, uptime and agility successes together? They are powered by real-time insights driven through IIoT-enabled connectivity.

$50B / year—the cost of unplanned downtime

The Challenges of Real-Time Access and Control of Big Data

Implementing an IIoT solution that can access and control Big Data in real-time for immediate business benefits must address the following challenges:

Difficulty assessing the volume and variety of IoT data:

Many factories utilize both modern and legacy assets and devices from multiple vendors, with various protocols and data formats. Although the controllers and devices may be connected to an OT system, they are not usually connected in a way that they can easily share the data with IT systems as well. In order to enable Connected Manufacturing and emerging IIoT use cases, manufacturers need a solution that can handle all types of diverse data structures and schemas from the edge, normalize the data, and then share it with any type of data consumer including Big Data applications.

Managing the complexity of real-time data:

In order to drive predictive analytics use cases, a data management platform needs to enable real-time analytics on streaming data. The platform also needs to effectively ingest, store, and process the streaming data in real-time or near-real time in order to instantly deliver insights and action.

Freeing data from independent silos:

Specialized processes (innovation platforms, QMS, MES, etc.) within the Manufacturing value chain reward disparate data sources and data management platforms that tailor to unique siloed solutions. These niche solutions limit enterprise value considering only a fraction of the insight cross-enterprise data can offer, while dividing the business and limiting collaboration opportunities. The right platform must have the ability to ingest, store, manage, analyze and process streaming data from all points in the value chain, combine it with Data Historians, ERP, MES and QMS sources, and leverage it into actionable insights. These insights will deliver dashboards, reports and predictive analytics that drive high value manufacturing use cases.

Balancing the edge:

Understanding the right balance between data processing at the Edge and in the Cloud is a challenge and this is why the entire data lifecycle needs to be considered. There is a troubling trend in the industry as companies choose to focus on one or the other without realizing they can, and should, do both. Cloud computing has its benefits for long-term analysis and large-scale deployment but is limited by bandwidth and often collects vast amounts of data but only uses a small portion. The value of the Edge lies in taking action at the edge where it has the greatest impact and zero latency, and then sending the most valuable data to the Cloud for further high performance processing.

From reactive to predictive modeling capabilities:

Predictive modelling capabilities are key to delivering insights, and current platforms provide little to no modeling or machine learning capabilities used to predict and prevent disturbances before they impact operations. When building these models, the ability to leverage disparate data sources or data types promotes a robust and well-rounded model resulting in stronger prediction capabilities.

Why Litmus

IOT SIMPLIFIED—Our Industrial Edge Computing Platform is purpose-built to simplify the complex challenge of collecting, analyzing, and managing data from hundreds of diverse machines. We can help you embrace Industrial IoT to optimize efficiency, increase overall equipment effectiveness, and minimize costs

Why Cloudera

HYBRID AND MULTI-CLOUD—Run analytics on the cloud platforms. Easily and securely move data and metadata between on-premises file systems and cloud object stores.

ANALYTICS FROM EDGE TO AI—Apply real-time stream processing data warehousing, data science and iterative machine learning across shared data, securely, at scale on data anywhere.

SECURITY AND GOVERNANCE— Use a common securuity model, role and attribute-based accesss policies and sophisticated schema, lineage and provenance controls on any cloud.

100% Open—Open source, open compute, open storage, open architecture and ope nclouds. Open for developers, partners, and open for business. No lock-in. Ever.

Scalability:

Many companies struggle to scale IIoT solutions. They might be able to effectively collect data at a single plant but when it comes time to add 10, 20, 30 more plants—the project may stop in its tracks. Successful implementations provide the ability to connect to any number of assets, manage devices, and deploy applications and analytics across any number of sites from one central location. And the only way to do this is to develop affordable capability to store deep histories of real time high sample rate data and the leverage the elasticity of your compute power delivered through the cloud.

Connectivity at the Edge:

While the IIoT provides manufacturers with powerful tools to improve product quality, optimize processes and improve agility, most manufacturers do not have the expertise required to extract this value on their own. Challenges at the edge include encountering connectivity protocols designed by each sensor and equipment manufacturer, adding or even replacing assets to the edge platform, the need to process and perform basic analytics at the edge, and a great deal of white glove coding required to integrate with third-party applications.

Industrial Internet of Things (IIoT) Time Has Come

Connected Manufacturing is at a turning point as use cases explode alongside a real, measurable shift in data types. Today, Industry 4.0 delivers a multitude of data types from IoT sensors (semi structured data), business process sources, such as ERP, MES, SCM or data historians (structured data) and computer vision (unstructured data). Real-time and time series data is growing 50% faster than static data and streaming analytics is projected to grow at a 28% CAGR. Streaming data growth is driven by the demand for real-time insights and more importantly, autonomous decision-making. As a result, legacy data platforms and data ingestion methods are inadequate to address today's needs.

This turning point is driven by the operational and intellectual maturity of the IIoT ecosystem including hardware, sensor, edge computing, data lifecycle, and cloud. Each deployment is no longer a novel one-off adventure, but rather a proven endeavour delivering operational value for established use cases. Experienced partnerships such as Cloudera and Litmus offer customers complete Connected Manufacturing solutions they can implement for immediate value without disrupting current processes.

Litmus provides manufacturing customers with data connectivity powered by one platform to collect, normalize, analyze and take action on real-time data at the edge. Cloudera delivers a unified data lifecycle platform that offers the ability to access, analyze and store structured or unstructured data on-prem, in the cloud or hybrid cloud solutions enabling machine learning and business intelligence solutions with consistent security, and governance. Together, the Litmus and Cloudera offering satisfies key manufacturing drivers from edge to AI and simplifies the deployment process so companies can immediately add value with IIoT.

Enabling the Connected Manufacturing Data Lifecycle

Litmus and Cloudera have partnered to offer an ecosystem for a complete, end-to-end solution that can be implemented securely to overcome these challenges and realize the value of a connected factory. Connected Manufacturing use case implementation rests upon an architectural framework emphasizing Big Data ingestion and management, leading to machine learning and streaming analytics. The following illustration examines the Data Lifecycle for Connected Manufacturing in greater detail.

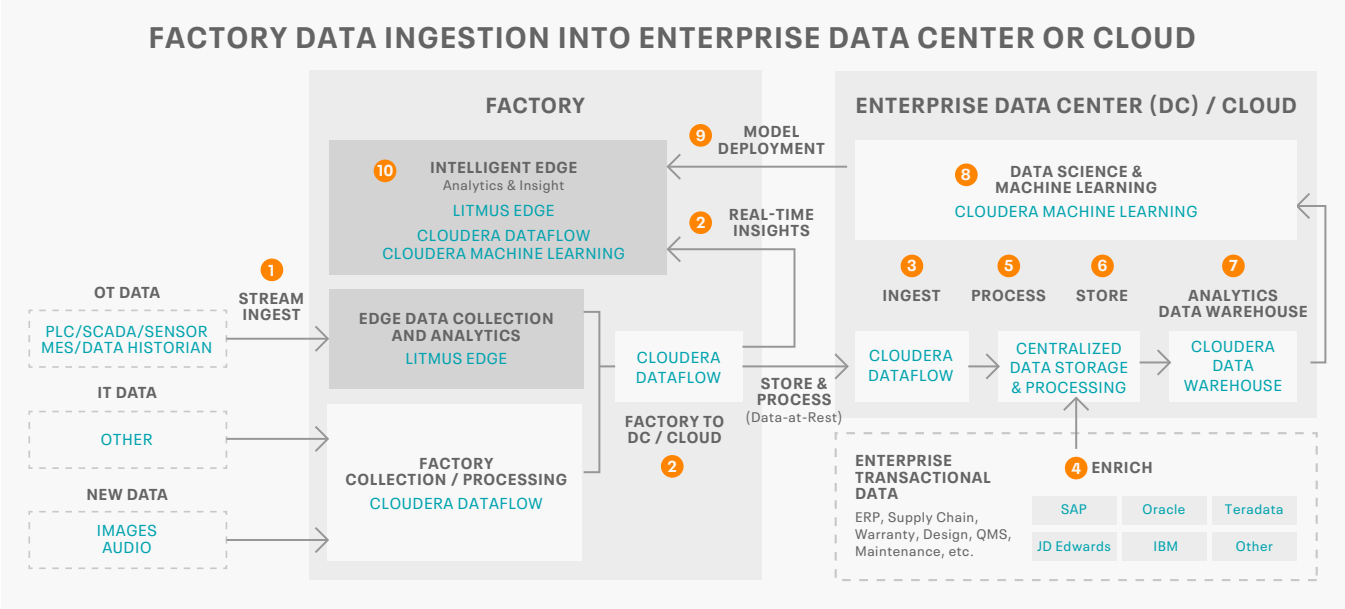

As outlined in the illustration above, the Data Lifecycle is comprised of the following steps:

- **Data Ingestion (Steps 1, 2 and 3):** The data ingestion phase begins with the connected devices themselves. Data from any number of sensors, PLCs, robotic systems, SCADAs and Data Historians with various protocols is ingested, normalized, and routed for downstream processing. Challenges addressed in this phase include the volume and variety of IoT data, reliable data flow control given often suboptimal network connections, and data normalization and analytics at the edge.

  Litmus specializes in enabling IIoT connections at the edge with the ability to support more than 250 industrial device protocols out-of-the-box with no programming and an intuitive interface. The complexity of setup, configuration and management is removed, and ready edge analytics and common KPIs and built-in third party integrations for rapid-time-to-value are added. The following value for Connected Manufacturing use cases is delivered:

  - **Purpose-built for industry:** Instead of offering generic IoT solutions, industrial markets are focused on to provide a ready-made solution that works from day one. The hard work is done with drivers, KPIs, and integration to cloud and business applications
  - **Secure and easy-to-use:** Litmus does all of the hard work so the plant manager can simply install Litmus Edge, connect a variety of machines with no coding, and then start collecting data in minutes. The system is simple and has security built-in at every step.

- **Integrated edge platform**: One of the most important considerations in an edge computing platform is easy integration with cloud and business applications. Litmus Edge integrates with Cloudera DataFlow, a comprehensive solution that ingests IoT data, log files, and data from various enterprise software systems, allowing users to easily leverage data from multiple sources, including both streaming data from IoT and business process data. CDF provides the abilities to aggregate, compress and encrypt connected manufacturing data, prioritize transmission of data from edge to the Cloud or Data Center, buffer data in the event of network interruptions and track the provenance and lineage of streaming data, providing confidence in the origin and usage of data. In addition, Cloudera and Litmus connect with AWS, Azure, Splunk, SAP and more, allowing manufacturers to keep the technology they have and realize value immediately.

  Data then has two routes: data can be returned back to intelligent edge devices running independent autonomous use cases (like anomoly detection used in process optimization on predictive maintenance programs) embedded in hardened edge hardware or flow into an on-premises data center or cloud data lake for further enrichment and refinement producing deeper insights.

- **Data Processing (Step 4 and 5):** Various data processing workloads (i.e. data quality, data combining, etc.) can be performed directly within the data lake, preparing the data for downstream use cases. Centralized data storage and processing provides massively distributed storage and processing engines for large data sets, storing and processing of any kind of data including unstructured data, semi-structured data and structured data (i.e., transactional data from ERP, maintenance, SCM, CRM data, etc.) and provide the ability to execute a wide range of data processing workloads in an extremely high-performance manner. Centralized data storage and processing addresses challenges in data processing operations such as time series data time stamp alignment, data quality clean-up (i.e. missing values, interpolation, standardization, etc) allowing for data to be further processed and analyzed with confidence.
- **Data Storage (Step 6):** Once ingested, real-time data is then delivered to the Manufacturing Data Lake, where it can be stored alongside a wide range of other data sources (i.e. ERP, MES, Quality or CMMS information). Cloudera enables storage on-premises, in the cloud or hybrid solutions.
- **Self-Service Business Intelligence (Step 7):** Once the data is stored and prepared within the Data Lake, it can be analyzed using data discovery or KPI-based business intelligence (BI) capabilities. Cloudera Data Warehouse is an enterprise-grade, hybrid cloud solution designed for self-service analytics enabling organizations to share petabytes of data to drive analytics and BI with the security, governance, and availability that large enterprises demand.

Faurecia

Faurecia is a tier-one automotive supplier specializing in automotive seating, interior systems, and emission control technologies. Like many manufacturers, Faurecia wanted to maximize production uptime and reduce manufacturing defects. Faurecia created an enterprise data hub for IoT that brings together and analyzes data from a variety of sources, including thousands of machines and millions of sensors, to help drive predictive maintenance and improve product quality.

"The cost of plant downtime can be significant, depending on the minutes of stoppage and every minute counts. With predictive maintenance, we are identifying correlations and patterns, and can take action before a breakdown occurs, avoiding unplanned production stoppage."

Jose Gascon, Manager of Faurecia's Core System Team

- **Learning (Step 8):** Predictive models are trained to detect anomaly patterns within large data volumes and from many diverse streams. At this stage, the depth and breadth of enterprise data is directly correlated to the power of the machine learning algorithm and hence the ability to produce an enterprise-wide value chain solution. Efficient learning is based upon clean data, access to all data, no matter its form, and a clear understanding of the path from model production to deployment. Cloudera Machine Learning can help operators accelerate data science at scale to build, test, iterate, and deploy machine learning models into production by taking advantage of massively parallel compute and expanded data streams. Using Python, R, and Scala directly in the web browser, Cloudera Machine Learning delivers a powerful self-service experience to data scientists to develop and prototype new machine learning projects and easily deploy them to production.
- **Data Monitoring (Steps 9 and 10):** Predictive models are then deployed to the edge via Litmus, where incoming IoT data is monitored in real-time (streaming analytics) to detect and identify conditions specified by the predictive models (i.e. specific sensor values that predict impending failures or process anomalies).

Implementing Connected Manufacturing: Time-to-Value

Choosing a data ingestion platform is based on a number of customer requirements. In addition to features, ease-of-use and time-to-value is an essential consideration. Many manufacturers have delayed embracing a data ingestion platform because they are under the impression it will cause some growing pains and come at a high price in both time and hard costs. The Litmus platform can be implemented cost-effectively for customers and can be up and running in just days. The beauty of the ideal IIoT platform is it will tap into existing technology on the shop floor, add a very important layer of intelligence, and then send that data across the enterprise for fast results.

Once the Cloudera data lifecycle platform is built and in place, move through the following steps for rapid time-to-value:

Proof of Concept—More than just offering product trials, the team helps customers develop POCs (proof of concept) from Day 1. Litmus Edge connects to devices in just minutes with support for 250+ assets with no programming required while other solutions take months to get to a POC. Connect an asset to Litmus Edge, see data immediately and start building dashboards.

- On or offline deployment with simple setup and no programming
- Start collecting data from one or several assets
- Develop a proof of concept to solve a business problem

Implementation—Day 30—Just one month in, customers are better understanding the data, seeing trends and taking action on new insights. They are beginning to think about what other data points they want to collect to expand use cases. With ready analytics and KPIs they can run analytics at the edge to achieve metrics like OEE and Asset Condition Monitoring.

- Calculate uptime/downtime, OEE and set up alerts
- Act on data at the source, where it has the greatest value
- Fuel big data and enterprise apps with valuable machine data

Production—Day 60—By day 60, customers can realize the true value of IIoT by sharing information across the enterprise. The data engine is fed, and the continuous feedback loop can begin. Plant managers can integrate their data scientists and orchestrate new applications. Just two months in, customers gain a very strong understanding of the data and share it with teams across the organization to solve problems like unplanned downtime, bottlenecks, and more.

- Develop machine learning models
- Orchestrate new applications for edge and IoT devices
- Share data and break down silos between OT and IT

Connected Manufacturing Use Cases

Manufacturers can see immediate value from the predictive IoT analytics solution provided by Cloudera and Litmus for various use cases, including:

- **Equipment and Process Monitoring:** Understanding your process and equipment conditions in real-time allow for more advanced use cases that improve yield and quality, but also enable remote equipment and process monitoring, and improve worker safety by stationing workers in the most environmentally and economically advantageous control room locations.
- **Improved Manufacturing Quality and Yields:** By increasing production yields and reducing costs associated with scrap and rework, major improvements in profitability can be achieved. This promotes the objective of 'predictable yield' that allows organizations to orchestrate their businesses rather than constantly react to production issues.
- **Reduced Issue to Resolution Cycle Time:** When unexpected quality or yield issues surface, data and powerful analytics capabilities are immediately at hand, providing rapid resolution and the ability to return to full production as quickly as possible. In addition, the knowledge obtained through resolutions can be easily incorporated into learning models to predict this behavior in the future.
- **Predictive Maintenance:** Anomalies in the production line often lead to unplanned downtime and costly maintenance. IIoT platforms can enable predictive maintenance to eliminate the instantaneous cost of failure. The ideal platforms come with KPIs for asset utilization, uptime/downtime, and more. Customers can set up real-time alerts to take action at the data source based on predefined events. Use cases often start with reactive maintenance, move to predictive with more data, and eventually prescriptive maintenance when the solution advances enough to make changes so the machine runs longer and better.
- **OEE (Overall Equipment Effectiveness)** is a manufacturing best practice coined in the 1960s that measures productivity. OEE evaluates how effectively a manufacturing operation is utilized by comparing fully productive time to planned production time. OEE calculation is based on three factors: availability, performance, and quality. Although OEE is not a new term, it is a perfect KPI for IIoT because modern technology offers readily available, real-time data to fuel OEE measurements and then help companies determine how they can improve OEE over time.

Litmus and Cloudera Provide Actionable Connected Manufacturing Benefits

Companies are reaping the benefits of Connected Manufacturing in their manufacturing operations and even extending it into their post-sales service organizations. Below is a summary of a diverse set of use cases that highlight how some of our customers are utilizing Litmus and Cloudera to drive predictive maintenance.

INDUSTRY	USE CASE	CHALLENGE	SOLUTION AND VALUE
Automotive Robotics Manufacturer	Equipment Predictive Maintenance Asset Predictive Maintenance	Automotive assembly downtime caused by robotic assembly equipment was costly. The non-IoT enabled assembly robots lacked real time insight into performance.	A program entitled **"Zero Down Time"** was initiated as a robotics monitoring solution leveraging an advanced analytics platform aimed at gathering, storing and analyzing IoT sensor data. The company is reaching their goal of zero down-time through predictive maintenance.
Oil and Gas Exploration	Equipment Predictive Maintenance	Downtime can cost tens of thousands of dollars a day on lost production caused by unexpected machine breakdown.	IoT data streaming from more than 70 trillion sensor data points and robotic automation systems feeds analytics. Predictive maintenance has **enabled six years of maintained production and slashed capital costs 80%.**
Steel Tube Manufacturer	Process & Quality Monitoring Throughput Optimization	Legacy systems lack the flexibility to combine data from various sources, create analytics forecasts and predictive models.	Implemented a machine learning solution to save costs throughout the manufacturing process. **Higher product quality** results in greater market share and lower costs is the result of **lower manufacturing energy consumption**.
Building Material Manufacturer	Asset Condition Monitoring	Lacked visibility on the shop floor with a large variety of assets and devices and no way to troubleshoot issues or improve asset performance.	Implemented on various production lines to collect machine data and run analytics for condition monitoring and anomaly detection with data sent to the Cloud for further processing. **Rapid deployment across 20+ plants with 70 percent asset visibility,** real-time analytics and advanced condition monitoring.
Oilfield Services Provider	Equipment Predictive Maintenance	Needed a closed-loop system for process improvement and predictive analytics.	Compiled operational equipment efficiency metrics and used them to adjust manufacturing processes and create alerts for key maintenance team members. An estimated **10 percent reduction in costs** due to predictive analytics and visibility.
Global Automotive Manufacturer	Predictive Maintenance	Lacked an edge to cloud solution to collect data from varied industrial systems and use it to create a machine learning model for predictive maintenance.	Collected data from 25 different industrial systems including PLC, Robotics, CNC, video systems, and RFID systems, normalized the data into one standard format at the edge, and pushed the data to the Cloud. Analyzed, created dashboards, pushed to central data lake to create **centralized predictive analytics, then pushed back to the edge for predictive maintenance.**

About Cloudera

At Cloudera, we believe that data can make what is impossible today, possible tomorrow. We empower people to transform complex data into clear and actionable insights. Cloudera delivers an enterprise data cloud for any data, anywhere, from the Edge to AI. Powered by the relentless innovation of the open source community, Cloudera advances digital transformation for the world's largest enterprises.
Learn more at cloudera.com

About Litmus

Litmus enables out-of-the-box data collection, analytics, and management with an Intelligent Edge Computing Platform for IIoT. Litmus provides the solution to transform critical edge data into actionable intelligence that can power predictive maintenance, machine learning, and AI. Customers include 10+ Fortune 500 manufacturing companies, while partners like Cloudera, Siemens, HPE, and Intel expand the Company's path to market.

Conclusion

Connected Manufacturing is proven to deliver tried and true use cases that improve product quality, optimize processes and improve agility. The Cloudera/Litmus partnership focuses on making data ingestion easy, eliminating the need for in-house expertise for deployment at the edge; providing the ability to leverage both OT and IT data in structured and unstructured format, and enabling the use of on-prem, hybrid or cloud storage and analysis platforms. Cloudera and Litmus provide an end-to-end solution, extracting value quickly and with high return. Many customers have accepted the challenge and implemented a 60-day proof of concept to enable digital transformation with meaningful operational improvements from Edge to AI.

Visit Cloudera's website and learn more at cloudera.com/solutions/manufacturing
Visit Litmus's website and learn more at litmus.io

Cloudera, Inc. 395 Page Mill Road Palo Alto, CA 94306 USA cloudera.com