



# Turning data into actionable insights

Lenovo Infrastructure Solutions  
for The Data-Centered

intel.  
**XEON**  
PLATINUM

**Lenovo**

The smart quotes, jokes, and zen statements about data don't exactly do justice to how important data has become today. While businesses find their way around the cloud, IoT, Industry 4.0, and everything else, data is growing at an exponential rate, faster than ever, in all kinds of ways. The digitization of transactions and business processes is a big part of this. The explosion of smart devices, security cameras, sensors, and, of course, social networks are making data grow even faster as more and more data is created and processed outside the traditional data center.

Data has been likened to the new oil<sup>1</sup>—if you can harness the power of your data, you can make better, faster decisions and potentially gain a competitive advantage. The question is, how do you effectively deal with the data explosion?

However, before you expect to be able to turn data into actionable insights, you should address these fundamental questions:

- What is the data your business generates?
- Where is this data, and how frequently is it generated?
- What are the questions you are asking of your data?
- How quickly do you need to gain insights, to stay ahead of the competition?
- What are the use cases you are considering for implementing your data solution?

Wherever you are on your data journey, Lenovo Smarter Infrastructure solutions, powered by Intel® Xeon® Platinum processor can help you find answers to these questions and offer Smarter Data solutions to help you transform into an analytics-driven, Data-Centered organization.





## **Make your business more data-driven with Analytics and AI**

We all know that analyzing data helps us make more informed business decisions. But it is easier said than done—especially given the volume of data that exists and how quickly it is expanding. The good news? We're here to help.

We have put together this step-by-step guide to help you make your business more data-centered, incorporating analytics and AI to drive consistent, actionable business insights for your customers, your competitive landscape, even your internal processes!

While it can feel like a tall (and stressful) task, automating processes, increasing efficiency, and continually innovating helps you stay ahead of your competitors and seize opportunities in this new data economy. Don't just manage your data. Master it.



## End-to-end analytics strategy

Finding valuable business insights requires an end-to-end analytics strategy across your data pipeline. Your analytics strategy should include:

- Defining the business problem(s) you want to solve with analytics is the most important step because data can confess to anything when tortured long enough. What you ask, you get.
- Making sure you have the right team members, experts, and resources—both internal resources and external partners
- Obtaining the right data for specific business problems you want to address
- Securing smarter infrastructure that supports current and future data processing requirements
- Detailing the problem-solving approach for each business case
- Continuously improving your practices to achieve desired results


It is crucial to consider every stage of your data pipeline (Ingest, Prepare, Analyze, and Act) when forming the strategy — make adjustments and improvements where they will have the most impact.



### STAGE 1:

#### **Ingest:**

Getting to the data sources



### STAGE 2:

#### **Prepare:**

Bringing all that data together



### STAGE 3:

#### **Analyze:**

Making sense before you make change



### STAGE 4:

#### **Act:**

Finding the insights that drive smarter action

So, let us get started.



## **STAGE 1:**

### **Ingest: Getting to the sources**

As obvious as it is, the first step towards analyzing data is collecting it. That means extracting it from various sources, detecting what has changed, then managing and moving it, efficiently and securely, from the data sources into a system where it can be stored and analyzed—whether you want to analyze it at the edge, in the data center, or in the cloud.

There are many approaches to collecting raw data, based on the size, source, structure, and latency:

#### **Data from application events**

This data, such as log files or user events, is typically collected in a push model, where the application calls an application programming interface (API) to send the data to storage. It can reveal user trends and provide valuable business insights.

#### **Streaming data**


This data is generated as a continuous flow. It might be a continuous flow of small, asynchronous messages such as those gathered from edge sensors, or much larger, such as video streams ingested from security cameras or monitoring equipment. Depending on the magnitude and time-relevance of the data, it's important to decide where to best process it and how much of it needs to be moved around the infrastructure.

We see a trend toward more real-time analytics of streaming data in various industries ranging from manufacturing to healthcare, retail to financial services. These use cases demand higher compute performance at lower power and the ability to deliver this high performance in non-traditional locations outside the data center. Together with our partners, we deliver solutions that provide those capabilities.

#### **Batch data**

This data tends to consist of large numbers of files transferred and stored in bulk or relational or NoSQL databases. It can be located on-premises or on other cloud platforms, and ingesting it requires high aggregate bandwidth between the sources and the target.





## STAGE 2:

### Prepare: Bringing all that data together

Data is often complex and layered, and messy. It comes in all different structures, sizes, and speeds and has to be cleansed, normalized, and prepared before it can be stored in a way that works best for your workload.

It might help to set criteria for placing or tiering your data depending on the urgency and usefulness, so the most cost-efficient memory or storage technology is available for each tier.

The goal here is to have your data in a format and location (whether at the edge, in the data center, or the cloud) that allows it to be easily accessed and queried. Consequently, your storage and data management platforms, compute performance, memory capacity, and network bandwidth must be aligned so the data your application needs is as close and as accessible as possible.

Not having the data at the right time and place for your applications and programs is what you want to avoid because they could end up being unable to analyze the data, or worse, give you unreliable insights.



## STAGE 3:

### Analyze: Making sense before you make change

At the risk of stating the obvious, focus on the use cases that deliver the most value.

That means taking it easy, at least at first. There is no need to go straight for deep learning when lower-hanging fruits in simpler analytics methods have yet to be used to find your data's potential.

Take a deep breath and start small. Achieving small and early successes can help garner organizational support and momentum for the bigger goals. Let us reiterate. Organizational support and leadership backing is crucial for any successful analytics journey. One of the biggest reasons businesses with big data analytics proofs of concepts failed is that they aimed for the stars and didn't leave orbit. Starting small also allows you to understand the appetite for data-driven decisions and allows you to course-correct if needed before the bigger investments are made.

### Preparing your infrastructure for analytics

The most important part of preparing for larger analytics and AI workloads is making sure your storage and network capabilities can keep up. Another reason to start small and scale-up.


Let's take the example of the beginning of the space age. The first launches did not consider even the return journey. Today, we have seen multiple successful tests of reusable launch vehicles and any number of creature comforts added to the lives of astronauts, which were never crucial elements for the designers and engineers of the first rockets. Yet astronauts aboard the International Space Station now enjoy such comforts as musical instruments, games, fresh-brewed coffee and most recently, fresh-baked cookies.

Likewise, as you progress in your analytics journey, larger volumes of real-time data for analysis generally need scale-out rather than merely scale-up infrastructure to keep pace with performance. As data grows, so will your needs, business cases, and even the tools you may use in the future. Building your analytics infrastructure with the ability to keep pace with your data is crucial in the longer term.



If you are used to CPUs, you may have heard that you need to invest in GPUs to support your heavier workloads - it is not that black and white, however. Most traditional Big Data Analytics and Machine Learning (ML) workloads (like clustering, regression, and classification) still run best on CPUs. GPUs are designed to process more data in parallel, making them preferable for Deep Learning (DL) training and, in some cases, for AI inference.

And don't forget about memory. Many real-time analytics workloads rely on in-memory databases that can quickly run out of capacity. You want to employ servers with the kind of memory you need for the workloads you have and will have in the future.



## STAGE 4:

### **Act: Finding the insights that drive smarter action**

At last! You have made it to the pipeline stage you have been waiting for —where you can extract meaningful insights from the data you have worked so hard to integrate, store, clean, and process.

Here, you can use in-depth exploration and visualization to understand your results better and transform them into something you can use to make an impact.

It is best to begin with your existing infrastructure and business intelligence (BI) tools. You can subsequently evolve as your capabilities and resources expand. It is easier and the outcome more predictable that way - and data scientists and business analysts can help you make the most of your findings, as well.

Now comes the goal of helping your users make sense of the results. Data visualizations, dashboards, and reports can come in quite handy. And as your data capabilities grow, so will the diversity in your user base. Your software needs to be flexible enough to evolve and develop over time to meet changing needs, as well as any new analytics and AI use cases.

This way, the end of this journey comes back to the beginning: choosing the proper use case, the right team, and the right data to find even more value within.

## Conclusion

As your business grows, your data grows, and so do the opportunities you have to make the most of it. The approach we've outlined, combined with a holistic consideration of your data pipeline and your needs, can help you transform your data center into a Data-Centered business.

## Choose Lenovo to get started today

As you become more Data-Centered, you need to face the choices you may have to make. Developing the entire skill-set in-house and maintaining it or finding the right partner to help you along your journey. While it may make sense initially to develop the competency in-house, that comes with its own set of challenges. Attracting the right talent, nurturing it, and retaining it can be both expensive and difficult, especially considering the fact that dedicated talent pools can get frustrated if the business cases do not allow for their skill and career development. In our experience, having helped thousands of customers to get the most out of their data, we can offer you all the capabilities, sans the risk of attrition, high cost, and difficulty in adopting new technologies, while also getting you started quickly. A few facts to help you make an informed decision:

- We work with industry-leading data, analytics, and AI software partners to design, configure, and test their software on our infrastructure. These solutions are optimized to perform and to scale alongside your data and your business. And whenever possible, we even validate our solutions with our partner's certification programs.
- Lenovo's AI Innovation Centers help you get results faster through Proof of Concept testing, Consulting, AI Workshops, ISV Testing, and Benchmarking.
- Lenovo's Storage and Data Management portfolio brings you the capability to easily manage your data from the edge to the cloud safely, efficiently, and securely.
- All our ThinkSystem and ThinkAgile servers are high-performance systems, currently holding 242 world-record workload performance benchmarks as of April 1, 2021—far more than any other vendor<sup>2</sup>.
- Lenovo ThinkSystem servers are engineered for always-on productivity and are consistently ranked #1 in x86 server reliability for the seventh straight year<sup>3</sup>.
- Lenovo offers a full stack of services for data analytics -oriented platforms.
- And Lenovo TruScale can even help you deploy solutions in a “pay as you go” consumption model, giving you reduced risk and full flexibility in how you procure and operate your infrastructure – while maintaining on-premises security and control.

That's what makes Lenovo Smarter Infrastructure solutions, powered by Intel® an excellent choice for your analytics and AI partner. We can provide you with a complete, end-to-end portfolio of servers, storage, software, and services, and we won't up-sell you on products and services you don't need.

## Let's talk. Contact Lenovo to schedule your data workshop.

Our Lenovo Services team has a data practice and can help take you from where you are to where you want to be — and it all starts with a complimentary half-day or full-day, hands-on workshop. Contact Lenovo today!

## References

1. <https://www.forbes.com/sites/forbestechcouncil/2019/11/15/data-is-the-new-oil-and-thats-a-good-thing/?sh=6aea44fd7304> ;
2. <https://lenovopress.com/lp1145-lenovo-thinksystem-continues-to-lead-the-industry-in-performance>
3. <https://lenovopress.com/lp1117-itic-reliability-study>

[Learn More](#)

Lenovo

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo.

Intel and Intel Inside is a trademark of the Intel Corporation or its subsidiaries in the U.S. and/or other countries. Other company, product and service names may be trademarks or service marks of others.

© Lenovo 2021. All rights reserved.