software
design
solutions
an applied visions company

HOW TO USE EDGE COMPUTING
to put more power in your IoT system

HOW TO USE EDGE COMPUTING to put more power in your IoT system

Executive Summary

The global edge computing market is expected to grow from $10.6 billion in 2018 to $21 billion by 2023. More companies are turning to edge computing due to the benefits it brings to IoT systems. Specifically, IoT systems become more powerful and energy-efficient when edge computing and Fog-based machine learning (ML) are properly leveraged.

The right approach involves smart segmentation of data analysis—knowing what data to gather and where to process it. Hint: The answer is not always the cloud.

We'll take a closer look at the edge and the Fog, providing tips on how to determine when to leverage them to your advantage and what issues to look out for along the way. We will also look at industries leading edge-computing adoption and how they are adding more power and capabilities to IoT devices, systems, and networks.

Table of Contents

What is edge computing?

At the most basic level, edge computing refers to locating the processing of data closer to devices operating “at the edge” of the network, rather than up in the cloud or in a central application server. This minimizes the amount of data being sent back and forth to the cloud so the data can be processed, the devices can be controlled, and reports can be generated. More of these activities can be processed at the edge.

You may be asking: How is this different from a typical IoT architecture?

In a traditional IoT system, sensors collect data, such as temperatures and vibrations. An on-property gateway collects this information and sends it to the cloud for processing. The information is stored at the cloud level, where it is processed and analyzed with big-data analytics or machine learning.

In edge computing, the IoT device transfers data to a local edge device for processing—before any data goes to the cloud. Depending on the use case, some or all of the data is sent on to the cloud or corporate data center.

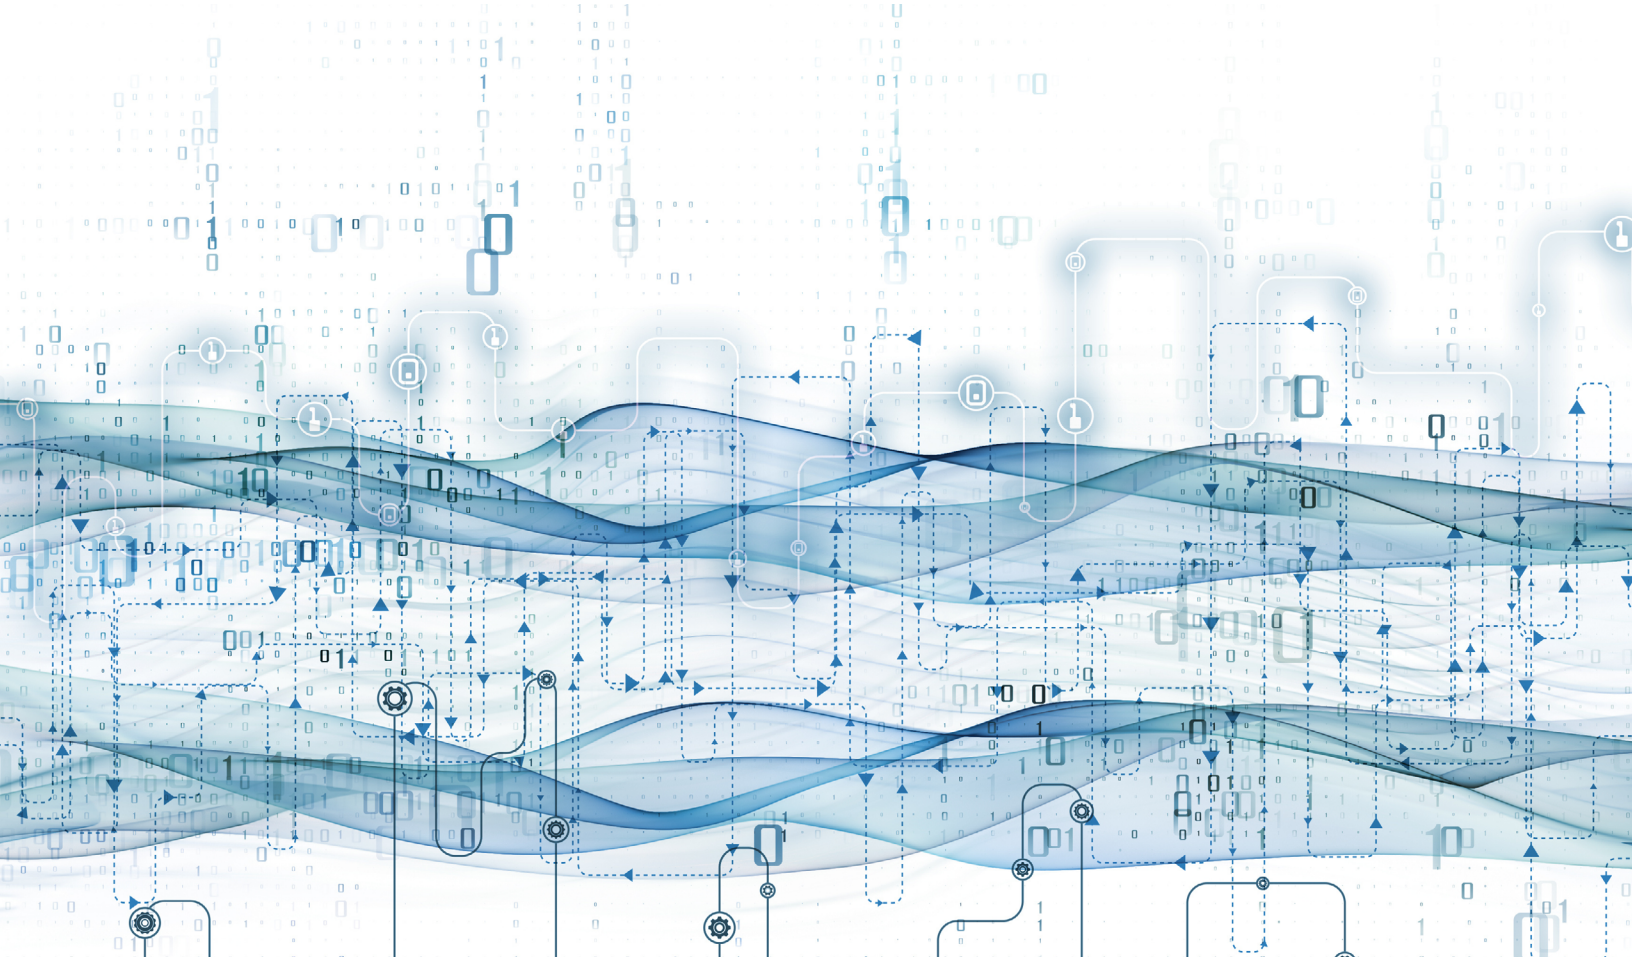

As edge computing takes hold, there are some key terms that are helpful to understand

Edge 

The place where the “users” (which may be machines or sensors) touch the network. In smart cars, the edge may be the automobile, while in Industrial IoT, the edge is likely a machine on the factory floor.

Edge Devices 

Any device (sensors, controllers, machines) that interacts with the physical world and produces or collects data.

Edge Gateway 

Think of this as the doorway between the edge of the network and the larger environment beyond it.

Fog 

This is a term that is becoming more common in edge computing discussions. The Fog refers to the space between the edge and the cloud. It encompasses edge computing, but also includes the network required to deliver processed data to the cloud or other central repository. You can get more value from data analytics when you leverage the Fog. More to come on this below.

Mobile Edge Computing 

This term refers to the use of edge computing to build and deploy 5G telecommunications networks. As major names in this space progress towards 5G, the term mobile edge computing is being used in more discussions.

Edge computing networks can connect to the cloud, but they do not need the cloud to operate. This may lead you to think the cloud loses some of its value with edge computing. That is actually not the case.

The cloud remains the ideal place for performing analysis that does not need to be done immediately or require historical information or a more global context. Predictive and preventive maintenance, asset health tracking, and the identification of patterns and trends should all still be done at the cloud level.

This information can be used to "teach" the local edge devices to improve local edge processing. In essence, the cloud and the edge work together to make IoT systems more powerful and robust.

A closer look at the edge versus the Fog

Any discussion on IoT should include a look at data analytics. After all, this is the real value of IoT—the ability to collect data from multiple devices or sources and to perform data analysis on that information to gain valuable and actionable insights. All so those devices can be controlled or better maintained.

To maximize the value of IoT data analytics, a distinction should be made between the Fog and the edge. The Fog is the area situated between the edge and the cloud. It operates close to the data source, but has enough power, bandwidth, and local context to make processing more capable than strict edge processing.

It's critically important to decide where to process the data in an IoT system. There are three tiers of data processing:

Local processing, which can be done in microseconds

Fog processing, which requires milliseconds

Cloud processing, which can be accomplished in seconds

The big questions are how much data you need to send up to the cloud and where to process the data within these three tiers. Many IoT systems (particularly remote Industrial IoT systems) lack the capability and connectivity to send all of the data collected to the cloud fast enough to be useful, making the Fog the ideal place for the heavy lifting of machine learning.

Many of us have also come to some realizations:

Cloud processing is not quite as cheap as we thought.

We don't really need all of the data up in the cloud.

If we need the data in real time or near-real time, sending it up to the cloud is not an option.

The solution is to segment data analysis and processing. Data analysis that is needed in real time can be done locally on the device or at the Fog level, depending on the exact time requirements and the processing capabilities of the sensors being used. Only some of this data needs to go up to the cloud, combined there with data from other systems and devices for more powerful and historical analysis.

Edge sensors are becoming more powerful. This is what makes local and Fog processing possible—but these sensors also cost more. The design of the system must balance functionality with cost. This is where an expert design team is important.

An autonomous car is an excellent example of maximizing data analytics at all three levels. Sensors installed on the vehicle have to be very smart and powerful, since all processing has to happen in real time locally (or at the Fog level) to keep passengers and pedestrians safe.

The car is connected to the cloud, but the car is not relying on cloud processing to make immediate decisions concerning how to drive and how to respond. The cloud is used for less time-sensitive issues, such as traffic updates and historical analysis that can improve how the vehicle operates overall.

What are the benefits of Fog-based ML and edge computing?

Connectivity 

With edge computing, IoT devices do not have to constantly be connected to the cloud. If you are deploying IoT devices in an area with poor connectivity, edge computing and Fog-based ML in particular are very useful.

Many Industrial IoT systems and "out in the field" systems fall into this category. Communication is limited due to terrain, bandwidth, or other conditions.

Data is processed locally and in the Fog; then, when communication with the cloud is available, the appropriate data is sent for further processing.

Reduced latency 

Latency is the amount of time it takes from when something is measured (such as a rise in temperature) to when it is recognized, processed, analyzed, and addressed. In edge computing, latency is measured in microseconds (locally) and milliseconds (in the Fog), rather than seconds (in the cloud). When processing occurs closer to the edge, data does not have to travel over a network for processing, increasing processing speeds.

There are many instances where speed of processing is critically important. The financial services industry is one area where large sums of money can be gained or lost without real-time data processing.

Autonomous cars (as previously mentioned) are another example. These vehicles produce an enormous amount of data that needs to be processed and analyzed instantly. There is no time to send all of the data to the cloud for processing. Edge computing allows the data to be processed locally, so the vehicle can respond quickly, keeping passengers and pedestrians safe.

Reduced network load 

As the number of IoT devices continues to grow, so does the volume of data. This exorbitant amount of data places a strain on the network, which slows down traffic and processing.

Edge computing reduces the load on the network by processing some of the data locally or at the Fog level. In some cases, only a portion of the data is then transferred to the larger network, providing additional benefits to network performance.

Protection from network outages 

Edge and Fog computing use the connection between individual sensors and a local data center. This reduces and sometimes eliminates the worry about cloud server downtime.

This is very important for critical industries such as chemical plants and government agencies. In the event of a local outage, data can be re-routed through other pathways, providing increased reliability.

Cost savings 

Cloud storage costs money. It is not always necessary to have a copy of all data sent to the cloud. Edge computing allows organizations to limit the amount of data sent to the cloud, reducing data management expenses.

Improved scalability and responsiveness 

Edge computing allows companies to expand at a reasonable cost through IoT devices and edge data centers. Companies can move into new markets without investing in heavy infrastructure.

Local edge data centers provide efficient, low-latency service. Businesses can be more nimble and agile.

Security

There is some debate as to whether security is a benefit or a requirement. We are including it in our list, but will discuss both sides of the argument.

Some claim edge computing is more secure, since less data travels over the network. Others argue it is less secure, touting that the edge devices themselves are more vulnerable to attack.

What all this means for you is that security should always be a main concern when deploying edge computing for an IoT system. Security should be addressed for both the edge devices and the larger network, including access control, data encryption, and virtual private networks.

Many hardware vendors are now baking security into the hardware, providing the building blocks for a secure system. Of course, your developers then need to take the next step and take advantage of these security features (such as Secure Boot) by turning them on, but at least the foundation will be there.

Which industries are the leaders in edge computing adoption?

Manufacturing 

IoT sensors collect a lot of data on the factory floor, from temperatures to airflow. Edge computing processes this data locally, so adjustments can be made immediately.

Local processing can provide fast-loop control back to the system to immediately address urgent issues, such as a cooling pump in an over-temp situation.

Temperatures can be adjusted, machines can be slowed when a worker approaches, and lighting can be adjusted for optimal visibility. These gains in safety and productivity are contributing to the adoption of edge computing across the industry.

Healthcare 

Local processing of patient data at the edge results in earlier detection of critical issues, increasing patient safety.

Staff can be alerted as soon as wheelchairs, beds, and other equipment are available, reducing wait times and increasing patient satisfaction. Inventories of medical supplies can be monitored to make sure critical supplies are ordered in a timely manner. Telehealth services can be offered in areas with poorer connectivity.

Retail 

Inventory calculations can be done with edge computing. Real-time analysis allows retailers to keep shelves stocked with high-demand items and assists with theft prevention.

There is no reason to send all of this shelf data up to the cloud. Only periodic inventory and product information needs to be sent to the cloud for storage and analysis.

Telecommunications 

Providers such as AT&T and Verizon are moving towards 5G networks, thanks to the low latency and massive broadband capabilities of edge computing. Verizon recently reported it cut latency in half with is 5G edge computing test, a testament to the benefits of this approach.

How can you implement edge computing technology in your business?

The first step is to define your data analysis needs to determine if edge computing is the right approach. Do you need real-time data processing? Do you have connectivity or bandwidth issues that edge computing can help resolve? Answers to these questions will help determine if an edge computing solution can benefit your business.

If the answer is yes, there are several steps in the process. There is a lot that goes into a successful edge implementation.

Your best bet is to work with a company that has expertise in edge and Fog computing so you know your implementation will be secure and deliver the maximum benefits to your business.

An expert team will build a Proof of Concept (typically within one to two months). A prototype will follow (usually within six months) that will deliver the functionality required. Future iterations will use a combination of off-the-shelf and custom hardware to deliver functionality within any cost limitations on the end product.

Below are some of the high-level issues you need to address in an edge deployment:

Design your network architecture

This includes your cloud servers, routers, access and edge nodes, gateways, and end devices.

Develop data governance policies

These should outline what data will be processed at the edge and what data will go to the cloud. Policies should also document who has access to the data. Back-up procedures for data recovery are critical and should be documented as well.

Identify the right team members

You will need team members responsible for transferring the data from the edge to the cloud, as well as people well-versed in data analytics.

Don't forget security

The security of your edge devices should be a primary concern. There are some companies providing security solutions to help secure microcontrollers at both the board and network level. Microsoft Azure Sphere is an example. If you take advantage of such a service, be sure you also give proper attention to physical security, as many edge devices are in public locations that may be difficult to secure.

Solutions such as Amazon AWS and Microsoft Azure provide products to help organizations design a successful platform and maximize edge computing benefits. The Open Fog Consortium is another resource for organizations interested in edge computing. It was formed to promote the adoption of edge and Fog computing across organizations and to develop standards for use.

Successful companies are taking IoT to the edge and reaping the benefits. When used properly, edge and Fog computing open new possibilities for IoT applications, systems, and networks. Knowing what data to gather, where to process it, and what truly needs to go to the cloud are major decisions that factor into the success or failure of your implementation. When you find the right balance between the edge, the Fog, and the cloud, the result is efficiency, cost savings, scalability, and reliability.

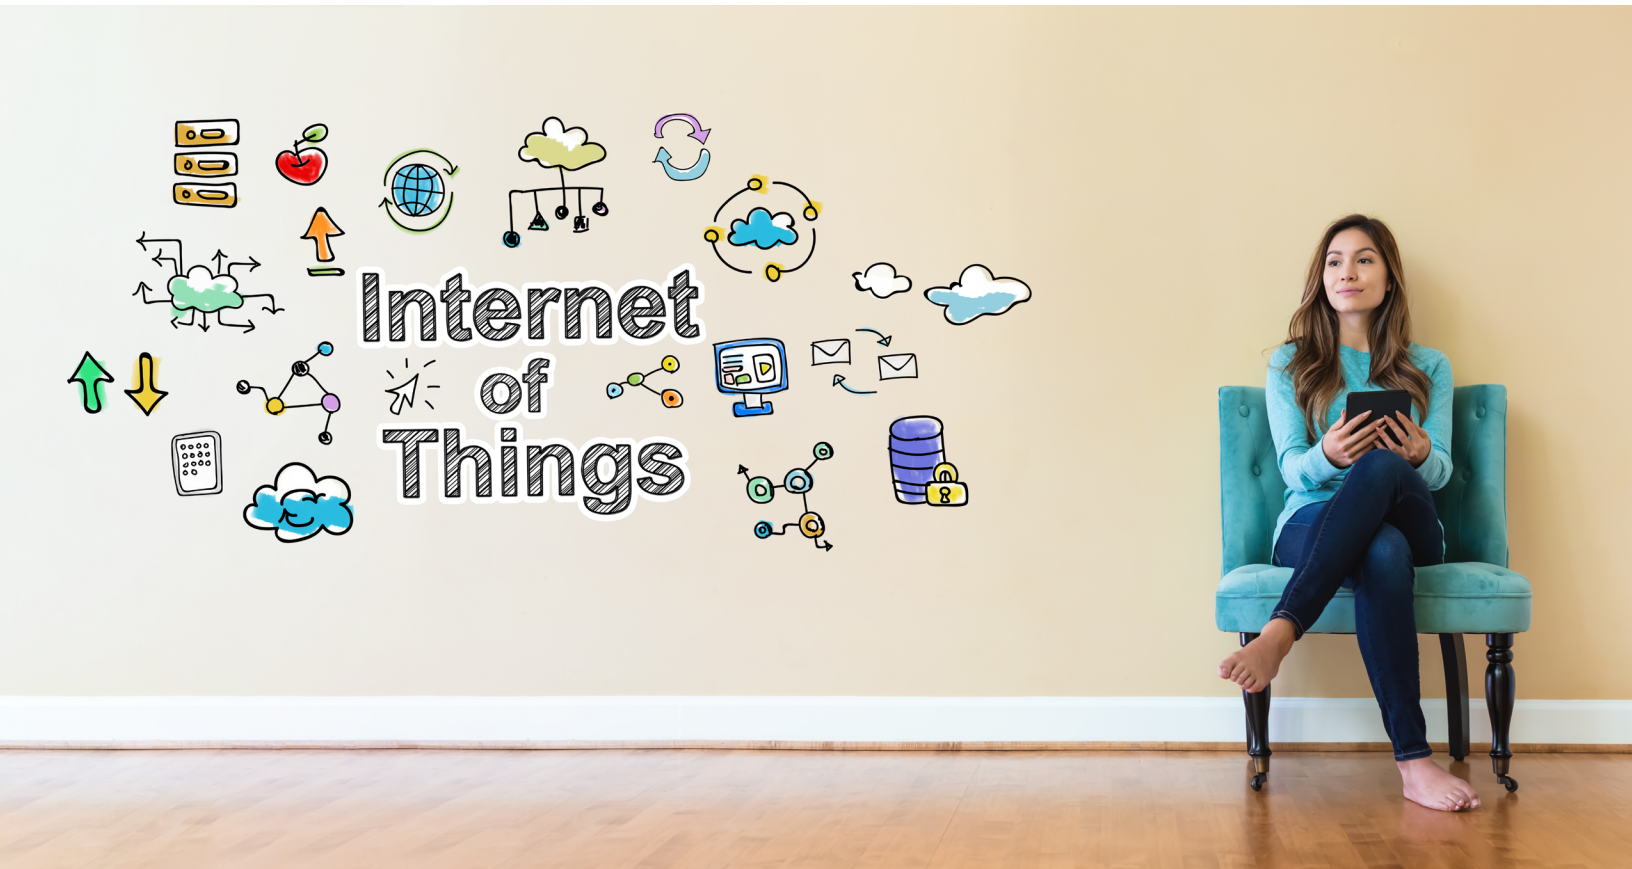

Software Design Solutions (SDS) provides embedded system and Internet of Things (IoT) software development, desktop application development, and software process improvement consulting. Software Design Solutions is also able to provide temporary software engineers for projects greater than the scope of its clients' abilities and to provide experienced programmers to companies who need to focus on other lines of business.

software design solutions

an applied visions company

Software Design Solutions, Inc.
4091 Saltsburg Rd. Suite S
Murrysville, PA 15668
412-282-8700
info@softwaredesignsolutions.com