Video Analytics Solution

ADLINK Media Cloud Server Assists in Intelligent Video Analytics

intel
IoT Solutions Alliance
Premier

ADLINK TECHNOLOGY INC.

www.adlinktech.com

ADLINK TECHNOLOGY INC.

Table of Contents

Introduction

Public security has become a focus of concern around the world. The wide availability of high bandwidth connectivity and increasing functionality of electronic devices has led to the recent advances in connected and Intelligent Video Surveillance (IVS) systems. Typically, IVS systems decode video streams from surveillance cameras, identify objects of interest using intelligent image recognition technology, analyze and track the objects of interest, automatically raise alerts when a security event is detected, and finally notify the user or relevant authority

An increasing number of IVS applications are appearing in different areas of society, including policing, finance, public administration, energy, transportation and education. China's IVS market is developing at a rapid pace, with increasing demands from various sectors and constant growth in infrastructure expenditure. At the same time, people's expectations for IVS equipment are also rising, for example, sophisticated behavior recognition, more precise alert functionality, better real time capabilities and more flexible deployment.

Beginning with Intelligent Video Analytics (IVA) network topologies and major components, this article then analyzes the current industry situation and the challenges faced by IVA developers. We then introduce ADLINK's Media Cloud Server which is specifically designed for processing video and big data, helping system vendors develop advanced IVA platforms.

IVA Applications

IVA platforms can be divided into two types: those deploying IVA hard-and-software at front end video capture points, and those with front-end camera systems soley responsible for image capture and transmission to back end servers for data analysis. Both types of IVA platforms have several pros and cons.

The main advantages of IVA platforms with front end analysis are a significant reduction in video data transmission and storage requirements. Video data is only sent to the back end servers when security event requiring an alert is detected. The majority of analytic processing is done at the front end, relieving the workload of back end servers. One disadvantage if intelligent cameras with integrated video analysis is a higher price tag. Furthermore, as each intelligent camera works independently, a given camera does not have access to the data acquired by adjacent cameras to track targets and analyze behavior in a broader geographic context. Front end IVA cameras also have less processing power than a back end IVA platform, and are therefore not able to implement more sophisticated analysis algorithms.

IVA platforms with all analysis performed at the back end put higher demands on network transmission and back end server processing capability. However, as network bandwidth becomes cheaper and cheaper, the performance and functional advantages provided by back end processing become more cost effective. Front end cameras with reduced functionality have an obvious cost advantages in large-scale deployments. Back end servers are also able to acquire video data from a wider geographic area and apply more sophisticated content analysis algorithms, resulting in lower miss and false alarm rates. IVA platforms with back end processing are also easier to manage and upgrade.

This remainder of this article will focus on IVA platforms with back end processing. The tasks performed by these platforms can be divided into three major parts:

● Front End Video Capturing

IThe video data acquired from surveillance cameras is compressed by encoding and transmitted to the back end system over a network via HTTP, RTP, RTSP or other protocol.

● Intelligent Video Analytics

Video streams from IP cameras are decoded into YUV format and pre-processed according to IVA requirements, including image enhancement, scaling, and color space conversion. Image enhancement can reveal faint, blurred or overly bright objects and make them available for analytics. Enhancement methods include infrared night vision enhancement, light and shadow suppression, object size filtering and image stabilization. The server then implements video analysis algorithms on the rendered image sequence.

Two approaches most commonly used in video analytics: one is to detect objects based on background modeling, which finds moving objects in comparatively static background images, categorizes the objects based on their shape and size, removes unrelated objects, compares the tracking trajectory against algorithm rules, and finally generates an alarm if criteria are met. This approach is often used for boundary defense, detection of objects left behind or abnormal loitering.

The second video analytics approach is based on feature recognition and is most commonly used for facial and vehicle license plate recognition. In contrast to moving object detection, feature recognition does not utilize background modeling, but rather focuses on the features of target objects. Therefore, analytic servers performing feature recognition need a database of feature descriptions in order to implement real-time analysis so as to find out similar objects. The richer feature database is, the more precise the results can be—however, the required computing capacity also gets greater.

● Back End System

IThe back end system receives video streams from the front end and processes the video to generate the analytics results, and displays the trajectory of target objects, statistical data and matching features on monitors, raising alerts if the results match pre-defined criteria. The results and associated video clips are also indexed for future retrieval.

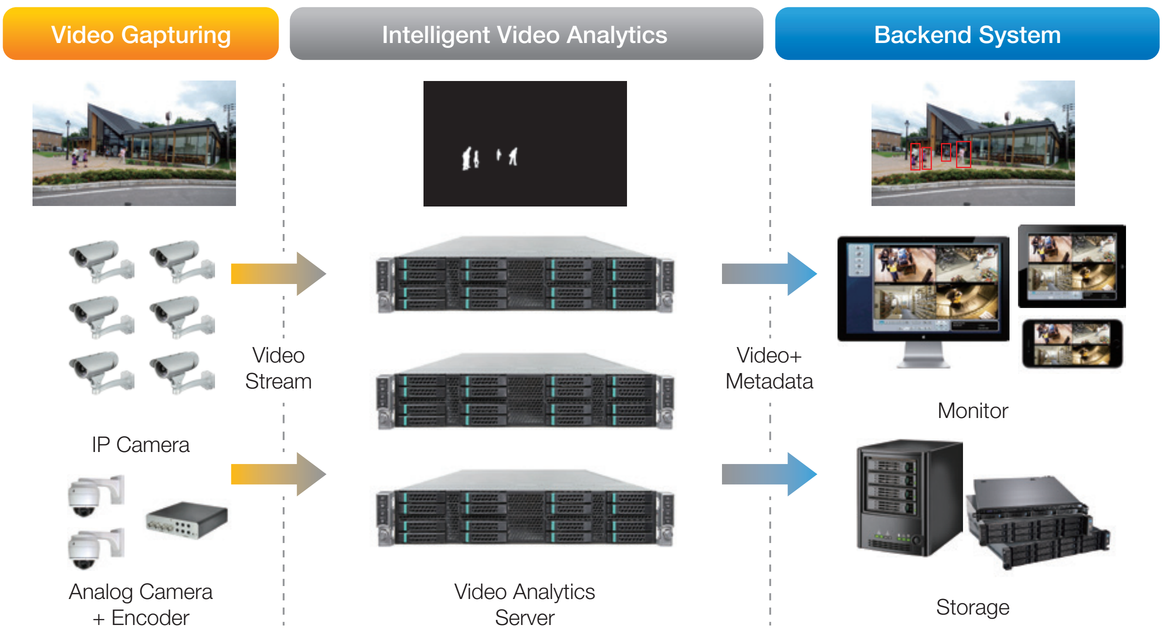

Figure 1: Network topology of an Intelligent Video Surveillance platform

Industry Trends and Challenges

Video surveillance creates a massive volume of data that is beyond human capacity to deal with. Control rooms operates 24 hours a day, 7 days a week, and security personnel can easily miss important events due to fatigue and/or slow reactions resulting from continually watching monitors for extended periods of time. Filtering out unimportant data and preserving significant information is a critical to achieving effective security. Intelligent video analytics, as an effective mechanism for raising alerts and alarms to prevent dangerous events before they occur, significantly raises the value and efficiency of video surveillance.

IVA is a combination of computer vision, image analysis and artificial intelligence technologies, aimed at creating relational mapping between monitored scenes and pre-defined events. IVA allows us not just to see what's going on, but to understand what has been seen. The relational mapping process involves video decoding, image processing and analysis algorithms. Provision and effective use of computer resources has become a key factor in implementing intelligent video analytics. IVA needs to compute and process huge volumes of data in order to:

(1) Decode video streams acquired from cameras. Resource consumption for decoding is directly related to the number of camera channels, resolution of data input and streaming format.

(2) Render decoded images before analysis, including image enhancement, scaling, and color space conversion. Computing resources consumed in this process depend on source image quality, objects and behaviors that need to be recognized, and the efficiency of analysis algorithms.

(3) Re-encode video for storage to facilitate future searches and replays.

The decoding and encoding necessary for IVA video management consume significant computing resources on their own, not to mention the resources required for image rendering tasks and intelligent video analytics. A 4th generation Intel® Xeon® E5 processor is only able to support two video channels of real-time transcoding of 1080p video from MPEG2 to H.264 using software en/decoding. Another important consideration for IVA platforms is how to effectively share the large volumes of decoded data with video analytics and re-encoding modules, as the video from a 10-second segment of video data [at 1080p H.264] decoded into YUV format consumes nearly 1G of storage space.

There are currently three major categories of IVA solutions on the market: digital signal processor (DSP), general purpose graphics processing unit (GPGPU) and general purpose CPU-based software transcoding. The following table summarizes and compares their characteristics.

(1) Currently DSP solutions are mainly used for front-end processing deployment. Usually a DSP processor is provided with optimal directives for enhanced graphic processing capabilities, such as convolution operations. However, compared with general-purpose CPUs, DSP is always disadvantageous in the number of computing cores and clock frequency. On the other hand, low commutability and flexibility among DSP solutions from different vendors result in high development cost.

IVA Solution	Processing Capability	Power Consumption	Flexibility	Development Cost
DSP (1)	Medium	Low	Low	High
GPGPU (2)	High	High	Low	Medium
CPU-based Software Transcoding (3)	Low	High	High	Low

Figure 2: Comparison of media hardware solutions

(2) IVA solutions based on general purpose GPUs (typically NVIDIA chipsets) adopt parallel architectures consisting of thousands of computing cores and can process massive amounts of data at high speed. However, in GPGPU architectures, the CPU and GPU share memory via a PCIe bus which becomes a bottleneck for data transfers and influences memory sharing efficiency. In addition, compared with general purpose CPUs, GPGPUs are usually more expensive and consume more power.

(3) CPU-based IVA solutions (typically Intel® Core™ or Xeon® processors) have advantages in development convenience and flexibility. However, general purpose CPUs are not optimized for video processing with most data processing tasks handled using software solutions, making it difficult to achieve high video data throughput even at high CPU usage.

ADLINK Media Cloud Server IVA Solution

IA solution than can simultaneously satisfy requirements for higher processing capability, lower power consumption and better flexibility will enable customers in tackling the challenges of video analytics in the cloud media era. The ADLINK MCS-2040, based on Intel® Media Software Development Kit (Intel® Media SDK) technology, is a dedicated media cloud processing server designed to meet the above mentioned requirements with added ability to handle big data. The ADLINK MCS-2040 comes in a compact 2U rack-mounted form factor and integrates 8 independent systems based on Intel® Xeon® processor E3 or Intel® Core™ i7 processors featuring integrated Intel® HD Graphics providing hardware accelerated graphics processing capabilities. Each individual system in the MCS-2040 can achieve 6 times the transcoding performance of pure software-based solutions running on comparable processors without integrated graphics while using only 1/8 of its processing capacity. The MCS-2040 also benefits from the usage of commercial Intel processors, providing the advantages of convenient debugging and smooth upgrading. In addition to the benefits provided by Intel's hardware developments and the Intel® Media SDK, ADLINK provides its MediaManager software package to make the MCS-2040 an Application Ready Intelligent Platform for video processing.

ADLINK MCS-2040 Media Cloud Server

The ADLINK MCS-2040 is a dedicated media cloud server designed for high-performance video processing. It allows users to use specialized encoding/decoding modules and programmable Execution Unit Array provided by the GPU in optimal ways. The encoding/decoding modules ensure fast processing of defined encoding/decoding tasks, while the programmable Execution Unit Array provides more flexibility by allowing users to run video analytics algorithms for much faster IVA processing.

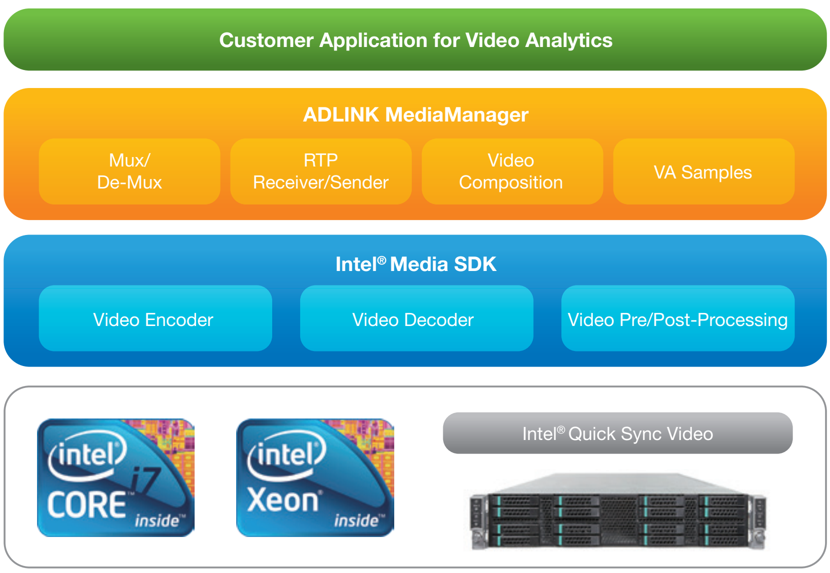

Figure 3: ADLINK MCS-2040 Media Cloud Server

The MCS-2040 supports up to 4 dual-system nodes for a total of 8 independent systems which support Intel® AMT technology for remotely system management. Each node features:

- Two independent systems (dual CPUs) which communicate via an onboard GbE interface
- Support for 4th generation Intel® Core™ i7/i5/i3 and Xeon® processors E3 v3
- Four DIMMs, up to 32GB DDR3 memory, 16GB per system
- Quad GbE RJ-45 per node, two per system
- Four hot-swappable 2.5" SATA drive bays, two per system
- 1600W redundant power supplie
- PCIex 16 slot reserved for 10G or Fiber Channel add-in card

Figure 4: MCS-2040's 2U rackmount design supports 4 dual-system media processing nodes

The MCS-2040 is the first hardware platform in the Intel® Media SDK ecosystem to be powered by an Intel Xeon® processor. It also supports GPU virtualization to allow GPU resources on the host to be shared among multiple concurrent VM clients. In addition to use as an IVA server, the MCS-2040 is also suitable for the following applications:

(1) High density real-time or off-line video transcoding server

(2) Multipoint Control Unit (MCU) for video conferencing, including traditional RTP/RTCP conferencing and Web Real-Time Communication (WebRTC) conferencing

(3) Cloud-based over-the-top (OTT) video services and applications such as remote virtual desktop, cloud gaming and virtualized set-top box

The following section summarizes the advantages of the ADLINK MCS-2040 with Intel® Media SDK-based MediaManager software solution compared to traditional DSP, GPGPU and CPU-based IVA solutions.

(1) Lower development cost and shorter development cycle compared to DSP solutions by allowing users to develop and test software in the native environment. Intel® Media SDK also ensures forward compatibility of software with future Intel® platforms.

(2) Better cost/performance ratio compared with GPGPU solutions by integrating the GPU in the CPU, eliminating the need for a discrete GPU card. This not only saves costs, but also reduces energy consumption during long-term operation. Recent trends among IVA platform solutions are to divide processing algorithms between CPU and GPU, thus requiring more efficient memory sharing and syncing mechanisms, which are provided by the MCS-2040's integrated CPU/GPU.

(3) Faster and more efficient video processing compared to CPU-based pure software solutions. The MCS-2040 provides GPU-based hardware acceleration graphics processing, which is especially beneficial for video analysis when all decoding, pre-processing tasks, and the majority of video analysis algorithms are executed by the GPU. The MCS-2040 capable of analyzing more channels at lower cost and power consumption than traditional CPU-based software solutions.

(4) Compared to generic Intel® Core™/Xeon® E3 platforms, the MCS-2040 is optimized for video processing applications and has up to 8 times the processing density, which is highly advantageous when deploying cloud-based data center solutions. In addition, the MCS-2040 provides side-band management functionality allowing for remote operations including system installation, startup/shutdown and reset.

Advantages and Benefits for Customers

ADLINK MCS-2040 Media Cloud Server can help users obtain improved efficiency in computer resource usage. It also provides a highly expandable and flexible system architecture which allows users operating cloud computing platform based on IVA technologies to realize higher efficiency and cost-effectiveness for reduced impact on the environment. This section summarizes the advantages and benefits of the ADLINK MCS-2040 Media Cloud Server for users in the areas of performance, availability, CPU/GPU memory sharing and big data processing.

● High Performance & High Density

ADLINK MCS-2040 Media Cloud Server combines the processing capability of 8 independent systems within a 2U machine, which is 4 times the density of traditional 2U dual processor servers. Each of the MCS-2040's 8 systems can support up to 40 channels of real-time 1080p H.264 decoding per system or 12 channels of real-time transcoding from 1080p MPEG2 to 1080p H.264 per system. Performing these en/decoding tasks uses only 20% of CPU resources, leaving ample reserve processing power for efficient execution of video analysis algorithms.

● ADLINK MediaManager Software

The MCS-2040 is bundled with ADLINK MediaManager software which builds on Intel® Media SDK's basic functions, providing capabilities including RTP/RTSP receiving and streaming and video composition to enable customers to build their IVA services. ADLINK MediaManager also includes sample reference code for realizing a proof-of-concept IVA architecture, which is illustrated in Figure 3 below. Customers can use this POC IVA to assess the performance of the MCS-2040 or use it as a basis to develop their own IVA software based on ADLINK MediaManager, thus reducing technical risks during product development and speeding up time-to-market.

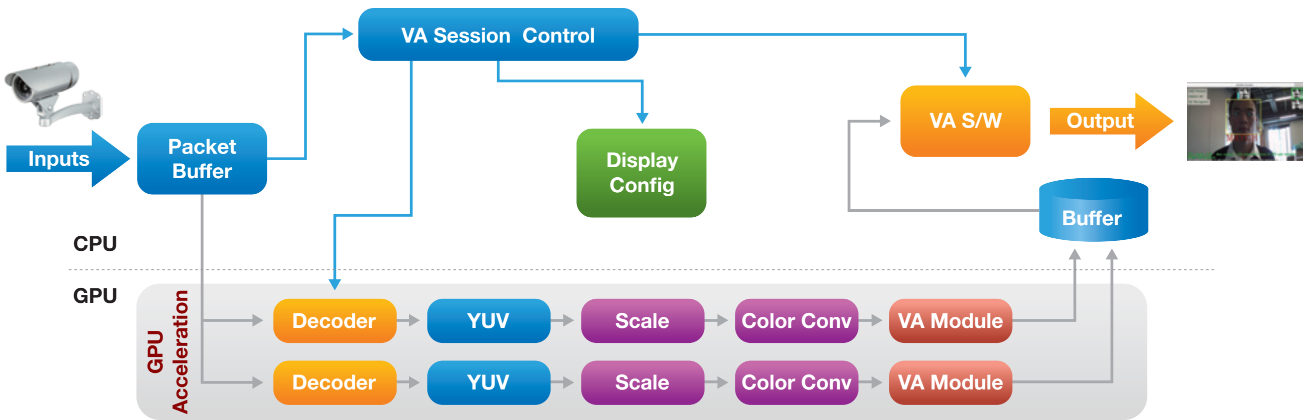

Figure 5: A reference IVA architecture with sample code provided with ADLINK MediaManager

High-efficiency CPU/GPU Memory Sharing

CPUs are good for logical operations and GPUs are superior at parallel programming. In order to maximize the advantages of both, a common practice in the IVA industry is to build a primary module for running analysis algorithms on the CPU, and move the tasks that are more suited to it to the GPU. Therefore, when the system is executing analysis algorithms, it is common for the CPU and GPU to share data. Avoiding copying data within memory and minimizing latency is one of the critical factors that determines IVA efficiency. ADLINK MediaManager helps to reduce data flow constriction in memory by taking advantage of OpenCL support to boost data sharing between GPU and CPU and therefore optimize memory efficiency.

Big Data Support and Cloud Deployment

The era of "big data" brings with it massive amounts of stored video data. The retrieval of valuable information by mining and analyzing big data is a main focus of the video surveillance industry. The ADLINK MCS-2040 uses the cluster computing approach to provide analysis capabilities many times more powerful than that of traditional servers. To provide better support for big data processing, the MCS-2040 has undergone a full range of tests using the most commonly used big data processing solutions such as Hadoop and Storm, ensuring that ADLINK customers can use the MCS-2040 to integrate these solutions more smoothly and conveniently.

To increase equipment usage efficiency, cloud deployments and virtualization technologies are often used in big data analysis. The MCS-2040 supports remote management functionalities including remote OS installation, remote system startup/shutdown/reset and remote KVM, allowing for convenient equipment operation and maintenance in the cloud. By supporting GPU virtualization, the MCS-2040 enables multiple virtual machines to share the processing power of an individual GPU, allowing IVA vendors to develop video analysis services on a subscription basis and provide specific analysis services for multiple subscribers on same host server. By moving virtual machines from one host to another during off-peak hours, operators can optimize the use of GPU resources to reduce equipment idle time and increase energy conservation.

Conclusion

In the era of "big data", more attention is being paid to Intelligent Video Analytics (IVA) technologies as traditional video surveillance is found to be too labor-intensive and have poor timeliness, low accuracy and low effectiveness. IVA relies on video algorithms to filter video content and extract meaningful information for subsequent action. Leveraging today's the powerful computing capacity provided by the processing technology, IVA can perform highly efficient analyses on massive amounts of video data and liberate surveillance personnel from repetitive daily tasks, allowing them to focus on potential risks or situations requiring attention. The market demand for IVA is expected to be high, with a need for accompanying IVA applications capable of processing huge volumes data. Customers are looking for solutions with more powerful computing capacity, lower power consumption, higher flexibility, low development cost and faster time-to-market. The ADLINK MCS-2040 Media Cloud Server is designed to meet these customer requirements with special consideration given to big data and cloud media processing. The highly efficient and energy-efficient MCS-2040 is a solid foundation for the new generation of IVA platforms.

Headquarters

ADLINK Technology, Inc.

9F., No.166, Jian 1st Rd., Zhonghe Dist.,
New Taipei City 235, Taiwan
新北市中和區建一路166號9樓
Tel: +886-2-8226-5877
Fax: +886-2-8226-5729
Email: service@adlinktech.com

Worldwide Offices

Ampro ADLINK Technology, Inc.

5215 Hellyer Avenue, #110, San Jose, CA 95138, USA
Tel: +1-408-360-0200
Toll Free: +1-800-966-5200 (USA only)
Fax: +1-408-360-0222
Email: info@adlinktech.com

ADLINK Technology Singapore Pte. Ltd.

84 Genting Lane #07-02A, Cityneon Design Centre,
Singapore 349584
Tel: +65-6844-2261
Fax: +65-6844-2263
Email: singapore@adlinktech.com

ADLINK Technology Singapore Pte. Ltd. (Indian Liaison Office)

#50-56, First Floor, Spearhead Towers
Margosa Main Road (between 16th/17th Cross),
Malleswaram, Bangalore - 560 055, India
Tel: +91-80-65605817, +91-80-42246107
Fax: +91-80-23464606
Email: india@adlinktech.com

ADLINK Technology Japan Corporation

〒101-0045 東京都千代田区神田鍛冶町3-7-4
神田374ビル4F
Tel: +81-3-4455-3722
Fax: +81-3-5209-6013
Email: japan@adlinktech.com

ADLINK Technology, Inc. (Korean Liaison Office)

137-881 서울시 서초구 서초대로 326, 802 (서초동, 모인터빌딩)
Tel: +82-2-2057-0565
Fax: +82-2-2057-0563
Email: korea@adlinktech.com

ADLINK Technology (China) Co., Ltd.

上海市浦东新区张江高科技园区芳春路300号 (201203)
Tel: +86-21-5132-8988
Fax: +86-21-5192-3588
Email: market@adlinktech.com

ADLINK Technology Beijing

北京市海淀区上地东路1号盈创动力大厦E座801室(100085)
Tel: +86-10-5885-8666
Fax: +86-10-5885-8626
Email: market@adlinktech.com

ADLINK Technology Shenzhen

深圳市南山区科技园南区高新南七道数字技术园
A1栋2楼C区 (518057)
Tel: +86-755-2643-4858
Fax: +86-755-2664-6353
Email: market@adlinktech.com

LiPPERT ADLINK Technology GmbH

Hans-Thoma-Strasse 11, D-68163, Mannheim, Germany
Tel: +49 621 43214-0
Fax: +49 621 43214-30
Email: emea@adlinktech.com

ADLINK Technology, Inc. (French Liaison Office)

6 allée de Londres, Immeuble Ceylan 91940 Les Ulis, France
Tel: +33 (0) 1 60 12 35 66
Fax: +33 (0) 1 60 12 35 66
Email: france@adlinktech.com

ADLINK Technology, Inc. (Israeli Liaison Office)

27 Maskit St., Corex Building PO Box 12777 Herzliya 4673300, Israeli
Tel: +972-54-632-5251
Fax: +972-77-208-0230
Email: israel@adlinktech.com

ADLINK TECHNOLOGY INC.

www.adlinktech.com