White paper

Secure remote services

Public cloud vs. private cloud (on-premise)

Find out more about

- Key differences
- Advantages for your company
- Risks in setup and operation

PHOENIX CONTACT

Introduction

The rise of public cloud solutions is causing many vendors of remote service solutions to shift their focus from on-premise solutions to cloud-based delivery models. This shift is also associated with a key question: "Which solution is the most suitable for my business?"
If you are wondering which option is more secure, accessible, and affordable, you will find the answers to your questions in this white paper.

According to a Gartner Study*, the cloud managed services landscape is becoming increasingly challenging and competitive.
Indeed, by 2022, up to 60% of companies will use the public cloud services of an external service provider, corresponding to a doubling of the percentage of companies from 2018 and a volume of 354.6 billion USD.
The advantages of the cloud are clear: It offers not only more agility and software that is always up-to-date, it also eliminates hardware constraints. Surprisingly, cloud vs. on-premise software continues to be hotly debated.

* Source: Gartner, November 2019
https://www.gartner.com/en/newsroom/press-releases/2019-11-13-gartner-forecasts-worldwide-public-cloud-revenue-to-grow-17-percent-in-2020

Essentially, the main difference between cloud and on-premise software is where it is installed. On-premise software is installed locally, on your business's computers and servers, whereas cloud software is hosted and operated on a vendor's server farm and mostly accessed through a web browser.

In addition to accessibility, there are of course other things that need to be considered when making a decision: software ownership, cost of ownership, software updates, backup strategy, and additional services such as security, support, and implementation. Let's explore all the pros and cons.

Content

1 Cloud basics

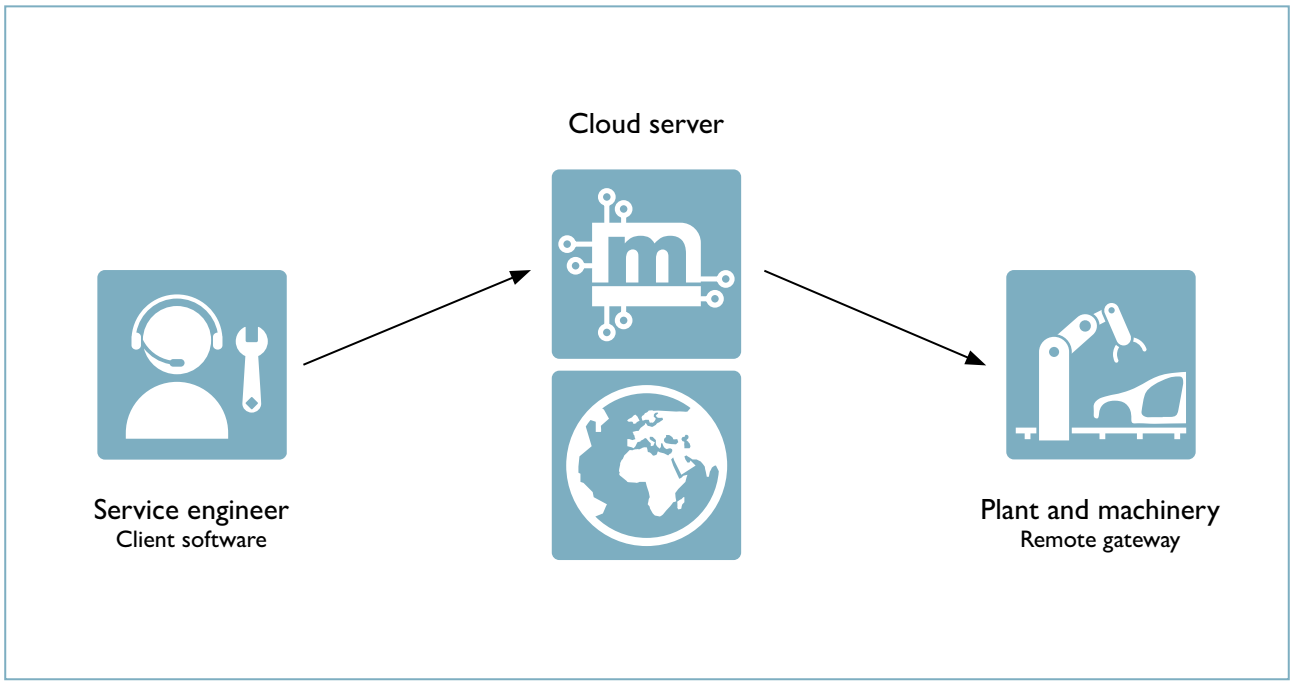

Typical remote services cloud topology

Cloud-based remote access is a new type of secure remote service enabling flexible remote access to machines in the field.

The associated network typology consists of three components:

1. Remote gateway
Remote gateways connect to equipment in the field to access and control it.

2. Cloud server
The cloud server is installed on a cloud-based platform such as Amazon Web Services or Microsoft Azure. It maps the connection requests and, after successful authentication on both sides, establishes a connection.

3. Client software
Client software is installed on the engineer's mobile device, PC, or desktop. The remote gateway and client software initiate outbound secure connection requests to the cloud server.

2 Cloud benefits

Cloud-based remote access solutions implement network topologies that enable the creation of outbound connections in the form of remote access tunnels. By doing so, they overcome the challenges presented by traditional VPN and remote desktop control technologies.

In addition, cloud-based remote access offers machine builders the following benefits:

User-friendliness

Plug and play remote access means technical configuration is no longer necessary. Security parameters, such as hash functions and encryption/decryption algorithms, are configured automatically. Machine builders do not need to configure these parameters; they just need to click on a button to establish a remote connection.

Virtual IP addresses make multi-point access effortless, with no field IP reconfiguration needed. Irrespective of the initial IP addresses set up by machine builders, cloud-based software assigns unique virtual IP addresses to machines. Machine builders can use these virtual IP addresses to establish several simultaneous remote connections. In addition, machine builders can use identical IP schemes for different field sites without worrying about address conflicts. This, in turn, cuts installation costs substantially.

Connections are centrally monitored and managed. The cloud server is the central point for establishing and managing remote connections. Administrators can monitor the traffic status and volume of each connection by connecting to the cloud server. Furthermore, administrators can easily manage client accounts, remote gateways, and certificates without having to reconfigure them regularly.

Enhanced security

End-to-end encryption between a remote PC and a piece of equipment prevents data leaks. The cloud server only routes traffic: It does not decrypt or store data that passes through.

Machine builders use remote access for troubleshooting, monitoring, maintenance, and diagnostics. Remote access is typically not required on a continuous basis, and can therefore be used on an as-needed basis. This helps to minimize security issues and reduce costs, especially when remote connectivity is based on a volume-dependent pricing option, such as with cellular technology.

Furthermore, machine operators take preference over machine builders in terms of remote access to all applications on their local network. Limiting access to only the applications machine builders need eliminates the risk of interfering with plant operations. Cloud-based access lets machine operators initiate or accept remote connections. Furthermore, machine operators can establish rules as to which services and applications machine builders can remotely use. They can also restrict access to specific sets of service engineers.

IT security policies are followed with no compromises. Cloud-based remote access solutions can build outbound connections using the IPsec VPN ports 4500 and 500, which means opening those ports for VPN traffic. Opening these

ports often means asking for trouble with IT and firewall managers. However, if the firewall-friendly service port 443 is used for this purpose (normally reserved for secure website access using SSL) or 80 (reserved for unsecured website access) is used for remote access, this does not present any issues for managing IT departments.
This solution can be used without hesitation, according to the IT security policies of the machine operator.

Flexibility and scalability

The client software is not limited to specific hardware platforms. Users can download client software to any mobile device, laptop, or PC and have remote access from anywhere and at any time, as long as they have an active client account.

Remote access to equipment lets users act as if they are locally connected. A transparent tunnel connects the client with the remote equipment as if they were on the same network. So, regardless of the remote equipment being accessed (PLC, HMI, etc.), and independent of the protocol used to pull data or used for programming, machine builders can remotely acquire data or program (remote) equipment. They can use their own software tools to do so, as if they were sitting next to the equipment.

Remote access simplifies network expansion by allowing network administrators to easily add and remove equipment and manage client accounts and certificates.

OEMs and machine builders require secure, easy-to-use, on-demand, and scalable remote access to their machines in the field.

Traditional on-premise access solutions are cumbersome and require IT/networkingknowledge, as well as changes in security/firewall policies.

Remote access backed by a cloud-based management infrastructure can provide the ease-of-use, flexibility, and scalability required by OEMs, without compromising on security.

3 Costs

A good way to look at the cost benefits of a cloud-based solution is by using the "iceberg" analogy – in other words, the bits you can see don't make up the whole of the picture. Let's start with the bits you can see:

- **Initial costs**
- **Subscriptions**
- **Software licenses**

With an on-premise solution, your setup costs are almost always going to be much higher. For that reason alone, if you don't want to make a significant investment, you're better off with a cloud solution – you simply pay a subscription fee each month/quarter/year, rather than "buying" the software.

This monthly subscription fee is almost always going to cost you more than paying the ongoing license fees for the software you have bought – so in that respect, an on-premise solution proves to be more cost-effective long-term.

It does not end there, however – on-premise solutions come with a huge number of hidden costs that don't show up on paper. Unless you have a very efficient way of handling them, you will normally end up paying significantly more with an on-premise solution than if you choose a hosted cloud solution.

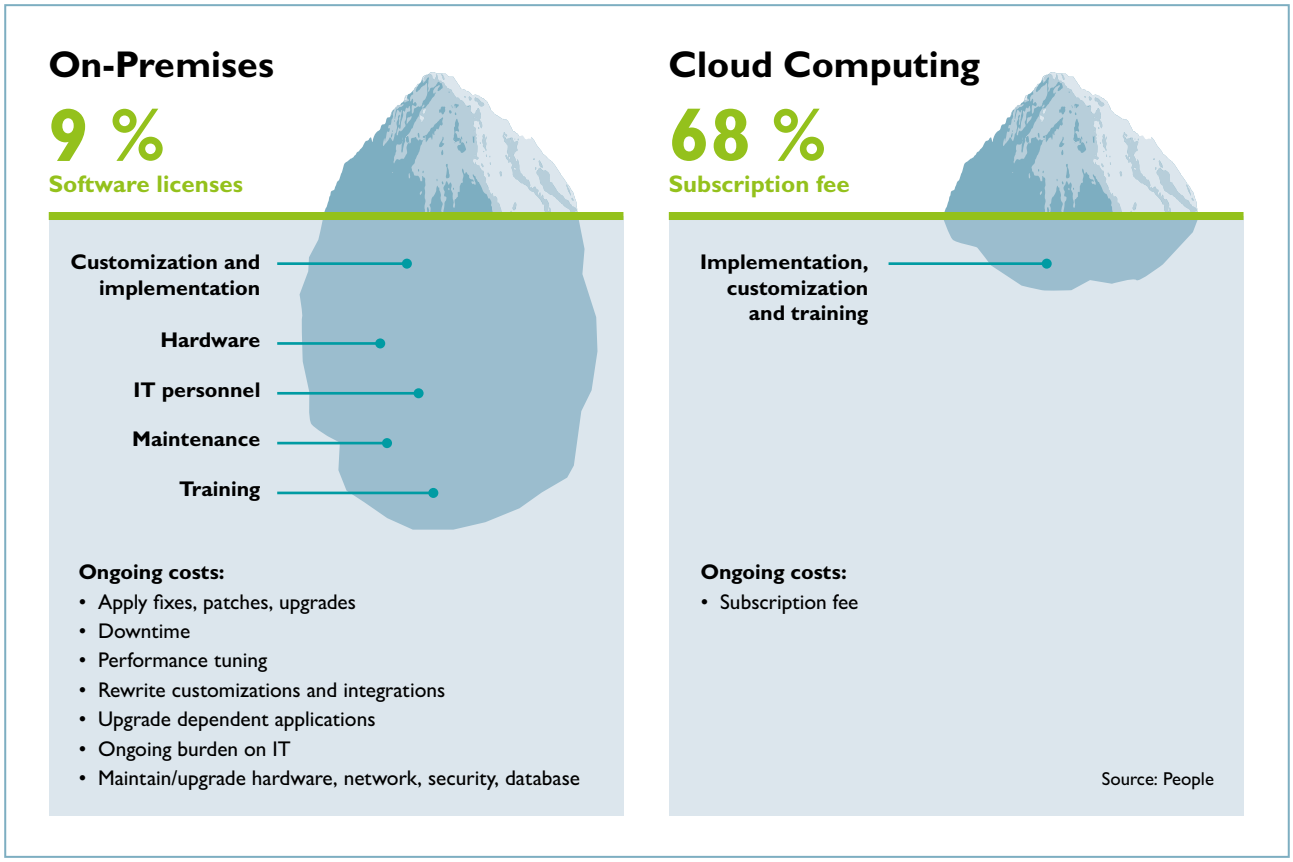

Cost advantages of cloud-based solutions: visible and hidden costs

4 Risks resulting from on-premise solutions

Building a private cloud in your on-premise data center can be a game changer. "Private cloud" implies the power of on-demand computing, at your disposal, with complete flexibility to construct a technical solution tailored to your specific application needs. A private cloud releases your dependence on the whims of providers like Amazon Web Services (AWS) and Microsoft Azure, allowing you to do things your way. For example, you can store data locally and manage compliance easily. Often, these methods also result in significant cost savings.

However, private clouds come with a unique set of challenges. Adopting a private cloud exposes your organization to several risks, some of which are not commonly known. What are these risks, and could they affect your decision to go for a private or public cloud? At the end of the day, the private cloud is still a cloud.

Risk #1: Security breaches

Private clouds can be less secure than public clouds. Public cloud providers have years of experience and top-notch expertise in security. In many cases they will have strategies, techniques, and tools to secure the various layers of the cloud stack. Certainly, public clouds are a bigger target for hacker attack. However, cloud vendors have an excellent understanding of cloud security concerns and how to mitigate them, which as a private enterprise you would have to learn.

Another concern is hybrid clouds. Security in a hybrid cloud is even more complex. When you shift workloads from the private to the public cloud there will be a transition from your internal security systems to those offered in the public cloud. In this transition, as traffic and apps are crossing from one system to another, there is a major risk of a "security lapse" that invites breaches.

Risk #2: Performance

Performance is a well-known concern in virtualized environments. Because of the highly dynamic nature of the environment, it is difficult to predict how changing loads at the infrastructure level will effect application performance and user experience.
Enterprises know their computing resources and how many machine instances they have in the public cloud, but there are other things that can impact performance – network bandwidth, latency and jitter, noisy neighbors on shared computing resources, access speed, and more.

The private cloud offers much more flexibility in how the cloud is built. You can select the hardware and software components, network infrastructure, and topology you believe will result in optimal performance for your use case. But will you really get the performance you think you will?

Just as public cloud vendors cannot always deliver the performance users need due to the complexities of virtualized and dynamically changing infrastructures, **you also will not always meet your performance goal with your private cloud.**

Hidden bottlenecks can occur in virtualized systems. Performance might vary depending on the current mix of workloads, software upgrades from VMware, OpenStack, or other elements of the system, and many other factors. Are you sure that your infrastructure will perform under all use cases and ambient conditions, including when it is upgraded?

Risk #3: Expertise and learning curve

Private clouds have been around for awhile, and many are built using VMware's software infrastructure, which is well-known and has a large user base. However, a growing number of private cloud projects are opting for the powerful and more cost-effective option represented by open source platforms. OpenStack is emerging as the new de facto standard for private clouds, but this platform represents a big unknown.

If you do not have accomplished OpenStack experts on your team – and there are not too many of those out there – it will be extremely challenging to get an OpenStack project off the ground. In the OpenStack User Survey 2016, users commented on the difficulty and complexity of working with OpenStack, although the platform is improving in maturity.

If you do not do things right in the early stages of an OpenStack deployment, you might experience significant difficulties later on. This might impact your ability to build the private cloud with the precise capabilities you need and to meet your timelines for milestones during your project.

Risk #4: Lack of visibility

One of the reasons to move from the public to the private cloud is to gain additional visibility into what's happening in the cloud. **A common perception is that once it is in your own data center, you will have much greater insight into things like workloads, usage, traffic, and performance.**

In the public cloud, there is no easy solution to gain insights on your network traffic at the packet level. Existing monitoring tools, like Amazon's CloudWatch and CloudTrail, do not let you "look inside the packets" to perform advanced diagnostics of network issues and prevent security problems.

In the private cloud, the situation is not much better. You'll face the problem of network traffic flowing between virtual machines (VMs) that does not touch a physical wire, and which is thus completely invisible to traditional monitoring tools. **This traffic can account for 80% or more of the traffic in a virtualized data center, creating a huge blind spot for IT teams.**

Risk #5: Limited scale

Many enterprises build a private cloud, as opposed to staying with a regular data center, to gain the power of on-demand computing and to be able to develop enterprise applications and services faster. However, at the end of the day, your private cloud's capacity will be constrained by your budget.

What happens if application usage is much higher than you expect? For example, if you run a customer-facing service and there is an "explosion" of usage, how will your cloud support it? You'll be exposed to the risk of overshooting your capacity, thus losing the economies of scale and cost savings that led you to build your private cloud in the first place.

The classic solution to this problem is a hybrid cloud, enabling "cloud bursting" from the private cloud to the public cloud if workloads overshoot your local resources. But setting up a hybrid cloud adds cost and complexity to your private cloud project. Furthermore, a common reason to build a private cloud in the first place is compliance with internal policies or external regulations.

There might be an internal policy stating that highly confidential data must be located on premise and must not leak to the public cloud or a legal requirement that data cannot leave the country. Using a hybrid cloud for peak loading in these scenarios might be problematic. How do you make sure only non-sensitive workloads are load-balanced between private and public clouds while keeping your confidential or regulation-affected information on premise?

Risk #6: Limited services

This applies to more than scale or cloud services and capabilities. On a public cloud like Amazon Web Services or Microsoft Azure, you have access to a plethora of cloud services, both native and third-party. They offer you everything from advanced management capabilities to auto scaling and high availability, storage services, instant provisioning of databases and huge data clusters, and so on.

You can use most of these capabilities in the private cloud, but you need to plan for them and then spend time and money integrating and deploying these features. In some cases, you even have to build capabilities from scratch.

In particular, high availability and resilience features provided by cloud services like Amazon AWS are difficult to recreate in-house. A feature like Multi-Availability Zones, which allows you to maintain replicas of machine instances in different data centers, is not be possible in most private clouds.

The bottom line is that **in a private cloud, you only have it if you have built it yourself.** If you didn't include a certain feature, functionality, or regular update in your project scope, you will be limited in your ability to innovate in the private cloud.

Risk #7: Data loss

According to data from Veritas*, many private cloud implementations are exposed to major risks of data loss. Data loss can occur on three layers: the hypervisor layer, the virtual machine layer, and the disaster recovery or backup system layer.

Due to the dynamic nature of a private cloud, traditional techniques for safeguarding data may not be enough and may not function in predictable ways across all scenarios. There are also numerous possible misconfiguration scenarios that can have disastrous consequences.

* https://www.veritas.com/information-center/enterprise-cloud-storage-ultimate-guide

Running multiple versions of VMware ESX, with some using virtual machine file system (VMFS) options unsupported by earlier versions, can lead to some VMs failing, data loss, and downtime.

If a critical application is running on two VMs, with one live copy and one backup copy, and one of them fails, there will usually be an automatic failover. If that failover instantiates the backup on the same physical host as the live copy, there is a single point of failure.

If there is a high-performance RAID 1 for production data and a lower performance RAID 5 for archiving and staging, there might be a mismatch. This will cause some VMs to write production data to the lower-performance storage, causing performance degradation or data loss.

5 Pros and cons of public cloud vs. private cloud (on-premise)

Cloud software advantages

Access anywhere and at any time
You can access your applications anywhere and at any time from any device over a web browser.

Affordable
No up-front costs are incurred for the cloud. Instead, you make regular payments, which makes it an operating expense (OpEx). While the monthly cost adds up over time, maintenance and support services are included, removing the need for annual contracts.

Predictable costs
Benefit from predictable monthly payments that cover software licenses, upgrades, support, and daily back-ups.

Worry-free IT
Because cloud software is hosted for you, you don't need to worry about the maintenance of your software or the hardware it is installed on; compatibility and upgrades are taken care of by the cloud service provider.

High level of security
Data centers employ security measures beyond what most businesses can afford. This means that your data is often safer in the cloud than on a server in your offices.

Quick deployment
Cloud-based software is deployed over the Internet in a matter of hours or days, as opposed to on-premise applications which need to be installed on a physical server and each PC or laptop.

Scalability
Cloud technologies provide greater flexibility since you only pay for your actual usage. You can also easily scale to meet demand, for example by adding and scaling back licenses.

Lower energy costs
When you move to the cloud, you no longer have to pay to power on-premise servers or to maintain their environment. This significantly reduces the amount you pay for your energy bills.

Cloud software disadvantages

Connectivity
Cloud solutions require reliable Internet access for you to remain productive.

Long-term costs
Although requiring a lower upfront investment, cloud applications can be more costly over the course of the system's life cycle, increasing the total cost of ownership (TCO).

Less customizable
Cloud software is typically configurable, but a cloud solution may not be able to cope with complex development projects depending on how it is hosted.

On-premise advantages

Total cost of ownership
Since you are only paying for your user licenses once, an on-premise solution can have lower total cost of ownership (TCO) than a cloud system.

Complete control
Your data, hardware, and software platforms are all yours. You decide on the configuration, the upgrades, and system changes.

Up-time
With on-premise systems, you do not rely on Internet connectivity or external factors to access your software.

On-premise disadvantages

Large capital expenditure
On-premise systems usually require a large upfront purchase, which means capital expenditure (CapEx) is often required. In addition, you need to factor in maintenance costs to ensure support and functionality upgrades.

Responsibility for maintenance
With an on-premise system, you are responsible for maintaining server hardware and software, security issues, data backups, storage, and disaster recovery. This can be an issue for smaller companies that have limited budgets and technical IT resources.

Longer implementation times
On-premise implementations take longer due to the time needed to complete installations on servers and each individual computer/laptop.

6 Summary

Why are cloud-based remote services better than on-premise solutions?

Cloud solutions are better than on-premise solutions due to more than just their flexibility, reliability, and security. They also eliminate the hassle of maintaining and updating your systems, allowing you to invest your time, money, and resources in implementing your core business strategies.

Providing real-time access to systems and data from a variety of devices regardless of the location and with guaranteed up-time of 99%, the cloud is becoming the number one choice for all enterprises using remote services.

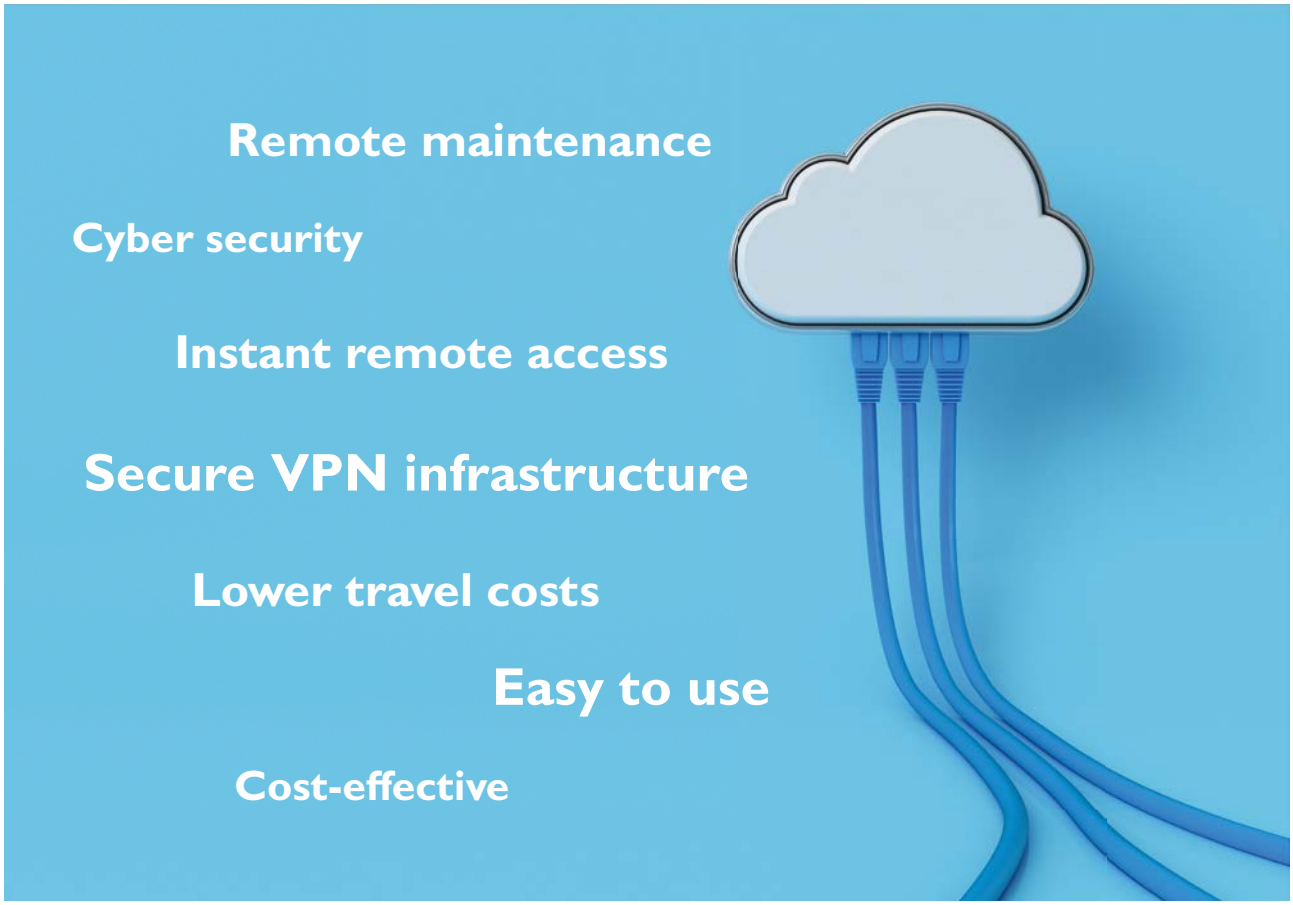

Contact

Secure remote services

As a strategic product manager, I am always looking for the best, not the most simple, answers for technology, design, user experience, and speed.
The results are sustainable and outstanding products that prepare the way for flexible access to your machines and systems in the field.

Find the best remote service solution for your company and make a consulting appointment.

https://phoe.co/mGuardSecureRemoteService

Markus Scheibenpflug
Strategic Product Manager
Communication Interfaces
Automation Infrastructure at
Phoenix Contact

mscheibenpflug@phoenixcontact.com

AI 03-21.006.L3