HYBRID CLOUD
TRENDS, TECHNOLOGIES AND TIPS FOR COMPANIES

LIFE IS FOR SHARING.

Contents

TREND TODAY, STANDARD TOMORROW
WHY THE HYBRID CLOUD IS BECOMING INCREASINGLY POPULAR

"The hybrid cloud is already the dominant modus operandi in German SMEs": this is the conclusion of a recent study by PAC[1] on behalf of Deutsche Telekom. According to the study, 85 percent of German SMEs in the manufacturing, retail and services sectors already use hybrid cloud models. However, more and more companies in other industries are also using hybrid cloud infrastructures. What are the reasons for this?

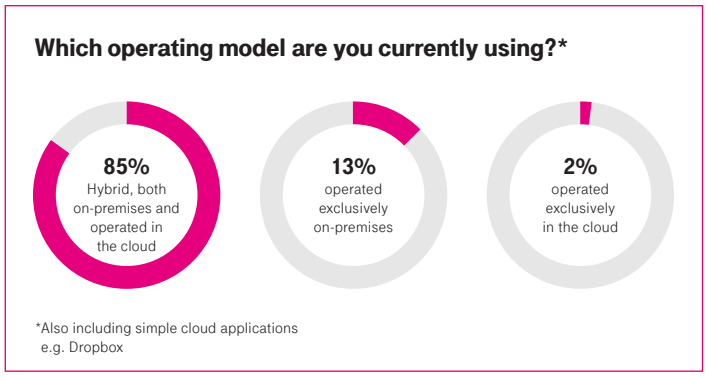

Best-of-breed principle

One is the best-of-breed principle: companies that use hybrid cloud models benefit from the best solution for their individual constellation. Business-critical data in particular can remain in their own data center to comply with internal and external compliance regulations. And for real-time applications, on-premises resources are also better suited than cloud capacity because companies benefit from the lowest possible latencies. For all processes that cannot be covered ad hoc by the company's own capacities, such as particularly computing-intensive processes, the public cloud is available at all times as a flexibly scalable infrastructure. Whether for machine learning processes, big data analyses or AI applications: With the public cloud, companies can fall back on practically limitlessly scalable resources that only have to be paid for as long as they are actually used.

Data gravitation

The applications themselves also contribute to the increasingly intensive use of hybrid cloud infrastructures: Software needs data to function. Because the use of software usually generates additional data, data tends to be collected where software is used – experts call this phenomenon data gravitation. Therefore, depending on the application, the time available and the size of the budget, it's necessary to weigh up whether it is more efficient to process data directly where it arises or whether it should first be transferred to another data center, for example to the public cloud. Using a single cloud topology for all processes is usually not cost effective. This is why more and more companies are opting for a sensible combination of public cloud and private cloud.

Network as bottleneck

In addition, while more and more companies are discovering the benefits of a shared IT infrastructure, they see the transfer of large amounts of data to remote data centers as an obstacle. According to a recent forecast by industry analyst IDC[2], the global volume of data generated and processed by companies, institutions and private individuals will increase fivefold to 175 zetabytes by 2025. This corresponds to a stack of DVDs with a length of 8.8 million kilometers – the equivalent of circumnavigating the earth 222 times. It is difficult to predict whether the capacity of the network connections will grow at a similar rate during this period. One thing is certain: Data should only be transmitted over networks if it can be done while keeping costs and latencies low and doesn't impair the transmission speed of important data packets in the network. For this reason alone, in the future companies will continue to process some data on-premises under certain conditions.

Legacy systems and legacy licenses

It's not always possible for some companies to freely choose their IT infrastructure. In some cases, companies are forced to run workloads on on-premises systems and/or in the private cloud due to historically grown IT systems, legacy systems developed in-house or legacy licenses. This is another reason why some companies are adding capacity from the public cloud to their existing systems rather than completely replacing their legacy applications.

The price makes the difference

One thing is clear: using a homogeneous cloud topology for every application doesn't pay off. Only those who distribute workloads sensibly to on-premises as well as public and private cloud systems benefit from the cloud principle through lower costs, greater flexibility and faster time-to-market.

5 TIPS FOR THE HYBRID CLOUD
WHAT COMPANIES SHOULD PAY ATTENTION TO

Public, private or hybrid? Not every cloud operating model is equally suited to every workload. Some divisions of a company may work with highly sensitive data, while others simply require maximum computing power, and others place the highest value on low latency. Hybrid scenarios whereby companies obtain IT infrastructure from both the public cloud and the private cloud are very much in vogue. This is confirmed by various studies, including a recent one carried out by market researcher PAC, which found that around 85 percent of German SMEs are already using hybrid operating models.

However, the architecture, set up and operation of a hybrid cloud infrastructure are highly complex. That's why it's important to have a detailed planning phase before introducing these kinds of solutions. What exactly needs to be considered? Here are five professional tips, using the Open Telekom Cloud Hybrid Solution as an example:

1. GET SUPPORT FOR SET UP, OPERATION AND SERVICE DIRECTLY FROM THE PROVIDER

When it comes to setting up and operating a hybrid cloud infrastructure, companies are usually on their own. That's because they don't typically receive support from the cloud provider for private instances. Companies only receive help from IT service providers from whom they purchase hardware components that are suitable for operating a private cloud instance and can be combined with certain public cloud instances. This is a complex undertaking with no guarantee of success, which requires several contractual partners at the same time – and makes troubleshooting more difficult in the event of a problem.

That's why users are better off with a provider like Telekom, which provides companies with its own experts for setting up and managing a hybrid cloud infrastructure. This ensures that the components are compatible with each other in the long term. And if problems or questions arise, the experts are also available after implementation. If you choose the Open Telekom Cloud Hybrid Solution you don't have to worry about implementation, operation, maintenance and service, because Deutsche Telekom experts will take care of all that while meeting the highest security standards. Deutsche Telekom is currently the only provider on the market to offer this kind of managed hybrid cloud model based on OpenStack.

2. ROUND-THE-CLOCK SUPPORT DIRECTLY FROM THE PROVIDER

The supplier should ensure that a competent contact person is available at all times to deal with any issues: Support also plays a decisive role in private instances. However, some companies opt to rely on their own personnel resources, which are usually not available around the clock. Or they may enlist the services of an external IT service provider. The ideal solution is when the provider offers support directly. However, most of them only support their own public instances if something goes wrong. As part of the Open Telekom Cloud Hybrid Solution, Telekom offers the same first, second and third-level support for private instances as for public cloud instances.

Which support offerings are very important to promote the use of hybrid cloud services in your company?

51 %
Consulting

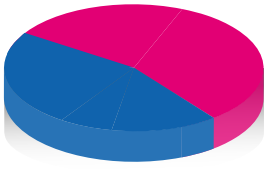

49 %
System integration services

According to the trend study "Hybrid Cloud in Germany" by PAC, around half of the companies would like support in planning and implementing hybrid cloud approaches.

3. OPTIONAL CONNECTION BETWEEN PRIVATE AND PUBLIC CLOUD

Anyone sourcing IT infrastructure from the hybrid cloud will usually want to use both operating models in parallel and divide workloads between public and hybrid as required. Some companies also want to isolate certain workloads in the private cloud, for example to comply with strict compliance guidelines. In the healthcare sector, for example, there are companies that have to operate certain workloads exclusively in private clouds without any connection to the shared infrastructure. With the Open Telekom Cloud Hybrid Solution, Deutsche Telekom is currently the only provider on the market that offers the option of operating private instances that are completely separate from the public infrastructure, but have the same look and feel.

4. ENSURE UNIFORM HARDWARE AND SOFTWARE BASIS

In order to rule out compatibility problems from the outset when introducing a hybrid cloud, companies should ensure that they choose a provider that builds and operates all instances on a uniform hardware and software basis. For many providers, it's the exact opposite. Hardly anyone offers exactly the same hardware that is found in the data center of a public cloud provider. This often leads to compatibility problems, for example at the network level.

With the Open Telekom Cloud Hybrid Solution, Telekom offers the option of operating private and public instances with identical hardware and software components. For example – just like in the public instance of the Open Telekom Cloud – the open cloud standard OpenStack is also used in the private environment. The hardware also corresponds to the same components as those used in the Open Telekom Cloud data centers in the state of Saxony-Anhalt.

Companies benefit from this in several ways: the compatibility of the instances with each other is guaranteed right from the start. If you develop applications in the private instance, you can also run them in the public instance without any problems. Furthermore, companies can more easily implement so-called bursting scenarios: In the event that certain processes require extremely high IT resources at short notice, resources can be spontaneously added from the public cloud – without the need for performing lengthy configuration. One example is in the area of high performance computing, where companies only need extreme capacities from time to time. Another is the retail trade, when double resources are needed in the web shop at short notice during the busy Christmas business.

5. TURN CAPEX INTO OPEX

A major advantage of Infrastructure-as-a-Service (IaaS) models is that they can be used as needed: renting the required infrastructure and thus transforming investment costs (CAPEX) into operating costs (OPEX). However, to benefit from this you have to ensure that all the components of the infrastructure – not just the virtual ones – can be rented when you set up a hybrid cloud. Anyone who buys the hardware for the private components of their hybrid cloud is taking the demand principle to an absurd level. If, on the other hand, you avoid the high one-off costs, you can instead invest in your core business and fully exploit the economic potential of the cloud. To do so, companies need a provider who can also supply the private share as an OPEX model.

The Open Telekom Cloud Hybrid Solution, for example, is based on this principle. Users choose the scope of the private part of their hybrid cloud solution, Telekom implements the desired configuration in the customer's data center or, alternatively, in one of T-Systems' highly secure data centers. Payments are made in monthly installments, the hardware is managed by T-Systems experts, kept up to date in agreed cycles and can be expanded at any time if required.

Which issues present a challenge for hybrid cloud use in your company?

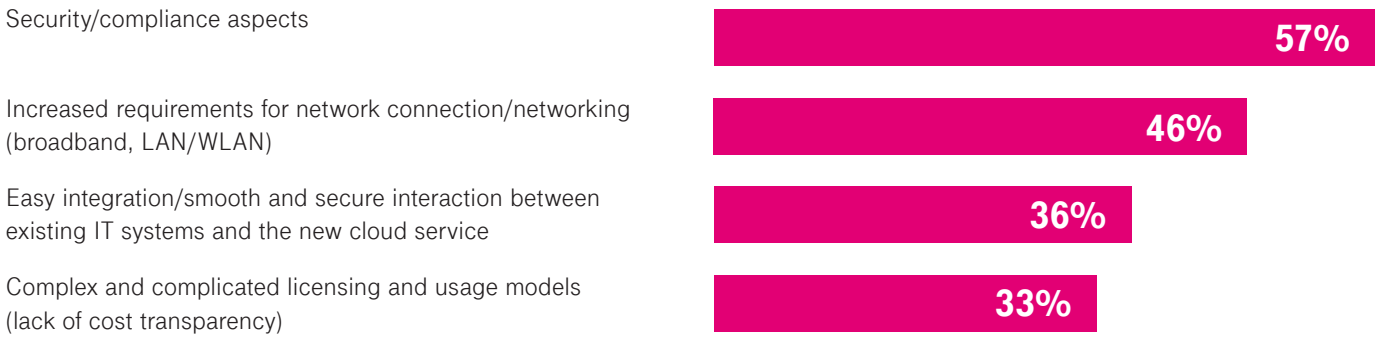

The major obstacles when it comes to a hybrid cloud solution in companies are in the security and compliance areas, followed by increased network requirements and the integration of existing IT systems.

Source: Trend study "Hybrid Cloud in Germany" by PAC

HYBRID CLOUD IN DETAIL
8 APPLICATION EXAMPLES

The hybrid cloud is gaining momentum: That's clear from, among other things, the current IDG survey[3] of 372 IT decision-makers from six European core regions. According to the survey, one in four companies currently has a hybrid cloud architecture – and that's expected to double by 2021. It's hardly surprising, since hybrid cloud solutions – when used correctly – combine the advantages of private and public cloud architectures. But what do concrete usage scenarios of a hybrid cloud infrastructure look like?

What type of provision or delivery models does your company currently use?
How will this look in three years' time?

Figures in percent. Multiple answers possible.

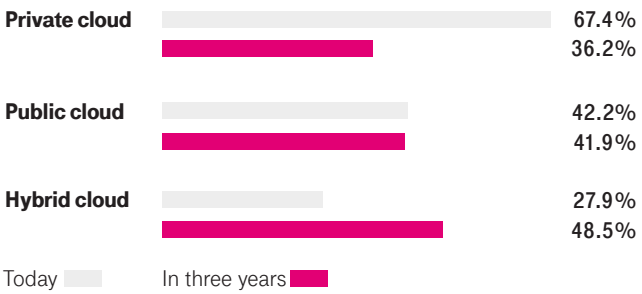

The hybrid cloud is becoming increasingly popular: the number of companies using a hybrid cloud infrastructure will almost double by 2021. (Source: Future Scape Predictions 2019)

1. SAFE HARBOR: PRIVATE ENVIRONMENT AS GATEWAY TO THE CLOUD WORLD

For companies, a hybrid cloud solution can serve as a gateway to the cloud world. The private cloud environment functions as a protected space for first steps and safe testing until the desired solution runs stably on cloud servers – in a second step it can be ported promptly into the public cloud if required. What is important here, however, is a hybrid cloud infrastructure with a uniform software and hardware basis to ensure maximum compatibility between the public and private clouds.

A suitable starting point could be legacy applications, for example: T-Systems' transformation projects show that around two thirds of business applications can basically be transferred to the cloud. This requires a thorough analysis and inventory of the existing IT landscape. T-Systems, for example, offers the so-called Cloudifier with standardized transformation services for entry and transition to the cloud.

2. BURST SCENARIO – THE PUBLIC CLOUD AS A BUFFER FOR PEAK LOADS

Whether it's the busy Christmas period in e-commerce, data-intensive simulations in product development or Big Data analyses in research: Many companies only need high-performance computing capacities from time to time rather than permanently. This is a case for the hybrid cloud, in which private and public clouds can be smoothly combined to form what is known as cloud bursting. This enables companies to flexibly absorb peak loads at any time by switching storage and computing resources on and off as required.

However, some companies need to store and process certain data in their own data center or in a private cloud. For example, to ensure they are protected against industrial espionage, to adhere to compliance regulations or to benefit from the lowest possible latencies. But the public cloud can also be an overflow basin for such companies: It depends on how such scenarios are implemented in detail. If business-critical data always remains in the private environment and only the compute service comes from the public cloud, cloud bursting can also be feasible for sensitive company data. This can be achieved either by anonymizing the data that is to be processed in the public cloud. Or by only providing applications in the public cloud with the exact information they need for a specific computing process. This makes the interaction secure and also saves bandwidth.

3. BACKUP AND DISASTER RECOVERY WITH THE HYBRID CLOUD

Hybrid cloud scenarios can also be used as backup or disaster recovery solutions. There are several possibilities for this kind of implementation – depending on the requirements of the respective company. On the one hand, the public cloud is an inexpensive long-term storage solution. For maximum security, for example, data can be stored in encrypted form in the Object Based Storage.

Companies that do not want to store certain data in the public cloud but still want to store it redundantly can also set up two private, separate availability zones (AZs) in which they mirror their systems. The distance between the AZs is important. They should be far enough apart so that both AZs don't fail simultaneously in the event of a possible disaster such as a flood or fire. But close enough to benefit from the lowest possible latencies. A benchmark that has become established among companies is a distance between the data centers of around 20 to 30 kilometers.

4. APPLICATION DEVELOPMENT WITH DEVOPS – DEVELOP ONCE, RUN ANYWHERE

Modern, agile application development increasingly integrates development, testing and operation. While the respective teams used to work independently and separately from each other, today they are linked by cooperation in a DevOps model. Hybrid cloud platforms enable teams to develop software faster and to shorten release procedures. Applications can be ported as needed between teams and their respective private or public environments.

What is important here is a hybrid cloud infrastructure that ensures a seamless transition between the public and private environments. The Open Telekom Cloud Hybrid Solution, for example, offers DevOps teams this kind of unified environment, which is based on the same hardware and software for public and private clouds. Each developed application can run in both the public and private environments – in line with the motto: develop once, run anywhere. And this is true even if the company in question doesn't work with container technology.

5. REAL-TIME DATA PROCESSING – MINIMUM LATENCY, MAXIMUM POWER

Edge Computing and the hybrid cloud cultivate a close relationship, because even with Edge Computing, companies can make use of decentralized computing and storage capacities as required. However, these are not located in a remote data center, but close to the action, at the edge of a network – hence the term "Edge Computing". Why is this necessary? One example is the real-time processing of data, where latencies can be kept as low as possible. For example, the transmission and processing of sensor data for autonomous driving practically require a data center at every intersection. The situation is similar for the control of industrial robots with AI algorithms, because low latencies also play a central role here.

There is a growing demand for IT resources that can handle processes with virtually no latency. Deutsche Telekom will soon meet this demand with their new Edge Cloud offering: mini data centers for real-time applications based on the Open Telecom Cloud technology, which can be installed and operated directly on the customers' premises if required.

6. IT DEPARTMENTS AS SERVICE PROVIDERS AND CLOUD BROKERS

The so-called shadow IT is flourishing – to the chagrin of many companies. According to Forrester Research, almost half of employees in companies now use technologies without the knowledge of their IT departments. Hybrid cloud solutions offer IT departments the opportunity to change things by building a unified service catalog and establishing themselves as the company's cloud brokers.

If the private and public parts are based on the same technology, the IT department will find it easier to build this type of unified service catalog. For applications with sensitive data that have to stay in the company, IT can then offer virtual machines with appropriate specifications. For less critical workloads that can operate in a public cloud, IT offers a VM with exactly the same specification – but at significantly lower cost. This means that users in specialist departments can select the services they need from private or public instances at the click of a mouse, without neglecting requirements such as scalability, security and governance.

7. PROVIDE REMOTE LOCATIONS WITHOUT NETWORK INFRASTRUCTURE WITH THE CLOUD

Whether it be an oil rig, space station or remote research facility: Not every location can be easily supplied with a fast Internet connection. With the Open Telekom Cloud Hybrid Solution, companies can also use IT resources in remote locations as a private cloud – regardless of the network connection. Telekom can implement the necessary servers directly where they are needed. A connection to the Internet or the public cloud is not absolutely necessary for operating them.

8. OUTLOOK: USING THERMAL ENERGY FROM DATA CENTERS SENSIBLY

Servers give off a lot of heat: To ensure an optimum operation, the hardware in data centers is usually cooled. This results in double the energy requirements – once for server operation and then again for cooling. But it's also possible to make sensible use of the thermal energy from servers. For example, supplying buildings with hot water and heating energy. In this way, companies not only save on the electricity for cooling the server cabinets, but also heating costs and they use their resources in a worthy and sustainable way. If the number of hybrid cloud architectures in companies doubles by 2021, more and more data centers will emerge in the future that are suitable for decentralized heat supply.

TELEKOM'S HYBRID CLOUD
UNIFORM SOLUTION FOR HYBRID, PRIVATE AND PUBLIC CLOUD

The Open Telekom Cloud is also available as a hybrid cloud solution – offering companies even more security, speed and convenience. Previously, the computing and storage capacities were available exclusively as a public cloud from Telekom's certified high-security data centers in Saxony-Anhalt. But now companies can also rely on these proven cloud resources based on OpenStack on top of dedicated servers, storage and network devices – that is, fully isolated infrastructure reserved solely for them. These can be located either at Deutsche Telekom's data centers, which companies access via a secure connection, or Deutsche Telekom can install the necessary hardware and software directly on-site at a customer's facility.

COMPLIANCE-CONFORM CLOUD USAGE

This means that IT resources based on the open standard OpenStack are also suitable for companies that cannot or do not want to use public cloud services for certain processes. For example, they wish to use the Open Telekom Cloud locally on dedicated servers at facilities abroad due to compliance or latency concerns. The certified Open Telekom Cloud already complies with Europe's General Data Protection Regulation (GDPR).

The Open Telekom Cloud Hybrid Solution makes it easy for companies to start using the cloud – even without prior cloud know-how. Experts from Deutsche Telekom handle all aspects of implementation, operation, maintenance and service according to the highest security standards. Another advantage: When required, companies can expand their dedicated resources with public cloud capacity from Deutsche Telekom's data centers at any time.

With the Open Telekom Cloud Hybrid Solution, companies can profit from the proven benefits of the Open Telekom Cloud with a dedicated solution as well. This gives companies access to a unique level of flexibility beyond any on-premises and public cloud limitations, incredible processing speed and Deutsche Telekom's well-known high standards regarding security and data protection.

The Open Telekom Hybrid Solution:
Either with the cloud based on OpenStack in the own data center ...

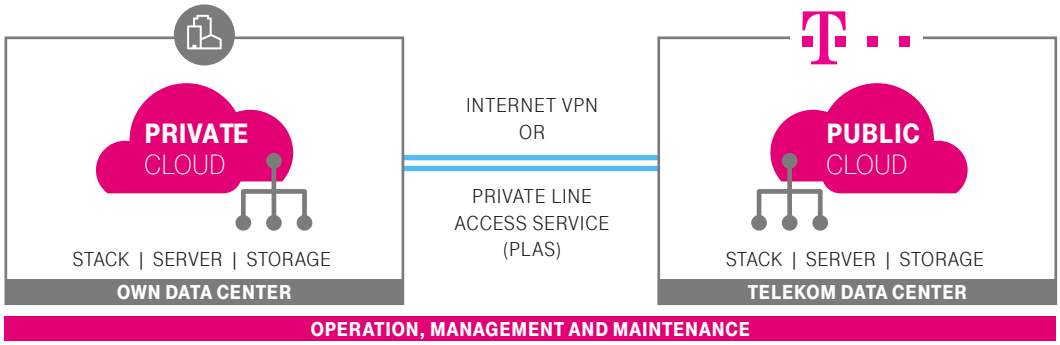

... or with a private cloud in Telekom's high-security German data center.

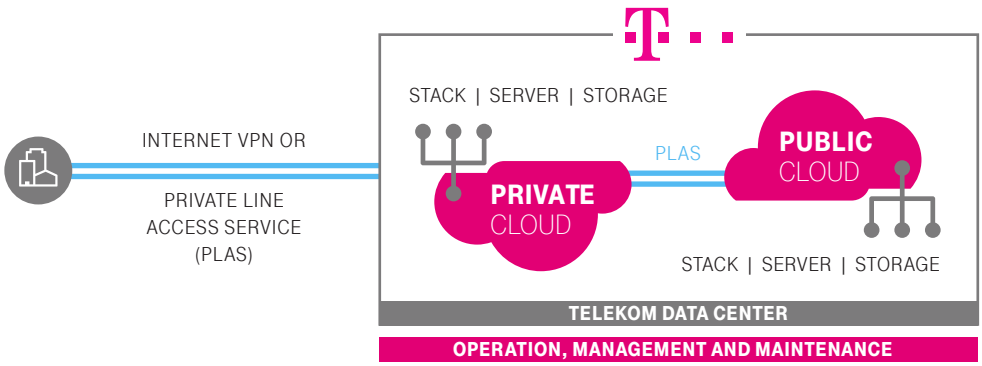

THE RIGHT DIMENSION
THE RIGHT HYBRID CLOUD FOR EVERY NEED

A lot of careful planning should go into every hybrid cloud infrastructure. Companies need to work out what kind of scope the IT resources should have. This isn't relevant for the public portion of a hybrid cloud, because it can be spontaneously scaled at any time. It's about the private part: how big should it be and what needs to be taken into account?

PACKAGE S

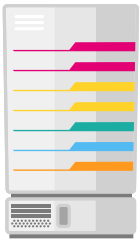

- 1 availability zone
- 176 physical cores
- ≈ 2.5 TB Ram
- 76 TB Block Storage
- Object Storage optional
- scales up to ≈ 1,000 physical cores
- SLA: 99 %

SCENARIO "SMALL": LOTS OF POWER WITH SHORT REACTION TIME

If the interlocking of resources is ensured, the next step is the dimensioning of the private stack. The size of the private part of a hybrid cloud infrastructure depends on the intended use. For example, some companies rely on the shortest possible response times for certain processes. If it's a question of milliseconds, then resources from the public cloud are not really suitable. That's why some companies like to add systems to their public cloud resources that work directly on-site. For example, to orchestrate production robots on the assembly line, to control autonomous transport vehicles in the warehouse or to ensure quality in production, where faulty components have to be detected and rejected within a fraction of a second.

PACKAGE M

- 1 availability zone
- ≈ 748 physical cores
- ≈ 11 TB RAM
- ≈ 76 TB Block Storage
- ≈ 112 TB Object Storage
- scales up to ≈ 4,400 physical cores
- SLA: 99.9 %

SCENARIO "MEDIUM": SUFFICIENT HYBRID CLOUD RESERVES

The "M" configuration of the Hybrid Solution offers significantly more computing power. With 748 physical cores, 11 terabytes of RAM and 112 terabytes of OBS, it provides sufficient reserves for particularly resource-intensive real-time applications such as those required in product development using virtual or augmented reality. For automotive manufacturers, working on virtual prototypes has been standard practice for years. Whether in design development, in determining the optimum drag coefficient or in chassis testing. High-performance cloud resources are particularly well suited to this.

Network · Compute Server · OpenStack Management · Block Storage · Object Based Storage (OBS)

PACKAGE L

- 2 availability zones
- ≈ 1,500 physical cores
- ≈ 22 TB RAM
- ≈ 164 TB Block Storage
- ≈ 240 TB Object Storage
- scales up to ≈ 8,000 physical cores
- SLA: 99.95 %

SCENARIO “LARGE”: MAXIMUM AVAILABILITY

The private instance of the Open Telekom Cloud Hybrid Solution in the “L” version offers twice the performance compared to size “M.” Another difference: Two availability zones raise the service level to 99.95 percent. This configuration enables companies to implement things like disaster recovery concepts. The resources offer the optimal prerequisites for this. The only requirement here is to maintain an adequate distance between the two availability zones. To achieve the best possible balance between availability and latency, a distance of around 20 kilometers between the zones is recommended, so that companies can continue to deliver even in the event of a disaster such as a fire in one of the two zones.

PACKAGE XL

- 2 availability zones
- ≈ 8,800 physical cores
- ≈ 128 TB RAM
- ≈ 980 TB Block Storage
- ≈ 240 TB Object Storage
- scales up to ≈ 22,200 physical cores
- SLA: 99.95 %

SCENARIO “EXTRA LARGE”: LARGEST POSSIBLE PRIVATE CLOUD CAPACITIES

The largest basic configuration of the Open Telekom Cloud Hybrid Solution is double that of size “L” and also offers two availability zones. There are no limits to the imagination here – just as there are no limits to dimensioning. The Open Telekom Cloud Hybrid Solution in size “XL” is highly scalable upwards. From AI development to big data analytics to deep learning – no matter what application scenarios companies want to implement, the “XL” Hybrid Solution makes it possible.

Network Compute Server OpenStack Management Block Storage Object Based Storage (OBS)

AT THE EDGE OF THE NETWORK
HOW EDGE COMPUTING AND PUBLIC CLOUD COMPLEMENT EACH OTHER

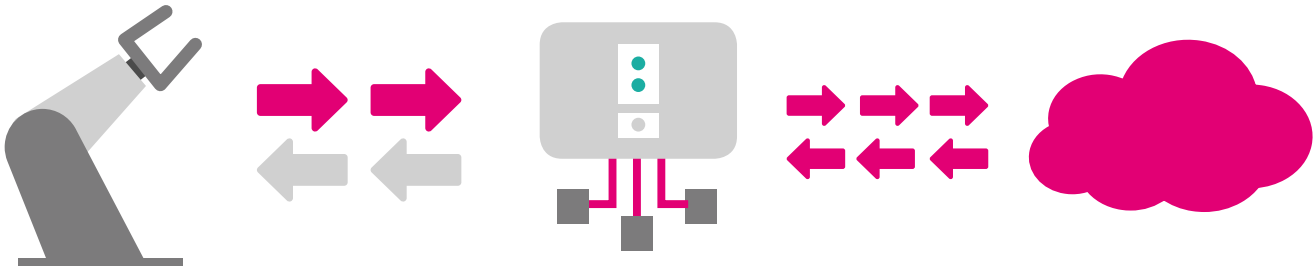

Edge computing is becoming the extended arm of the cloud: According to a recent forecast by market analyst firm IDC[4], in three years' time 40 percent of cloud installations will already have edge computing capabilities. It's a trend that Telekom is also currently observing: Data is the new oil. That's why it's only logical to bring refineries as close to the source as possible. But that doesn't mean that edge computing will replace cloud models. Edge nodes always need an instance to control them. If edge computing resources and the public cloud are intelligently combined to form an edge cloud, then the best of both worlds can be used.

COMBINATION OF AGILE EDGE UNITS AND SCALABLE CLOUD

From summer 2019, Deutsche Telekom will be offering the opportunity to combine lean edge devices with the scalable infrastructure of the Open Telekom Cloud in the form of the Open Telekom Cloud Edge. But why exactly do companies need cloud services if they can already process data locally with computing resources? Edge units usually consist of small, agile and lean hardware that takes on predefined tasks at the point of origin of the data, which often involve low latencies. In this tandem, the public cloud can take on more complex downstream tasks for which scalable computing resources are essential.

EDGE CLOUD EXAMPLE: AUTOMATED GUIDED VEHICLES

The most common examples include Internet of Things (IoT) applications, such as the control of Automated Guided Vehicles (AGVs) on a company's premises. There's a division of labor here: The vehicles transport pallets, for example between the warehouse and the production line. And an appropriately programmed edge node clears the way: it can open roller shutters, for example, if it detects that an AGVs is approaching the position and has to drive through it. Or prevent vehicles from colliding. This is done in real time to ensure a smooth operation and maximum safety. An edge computing system therefore serves as a control unit directly on-site.

PLANNING AND ANALYSIS WITH CLOUD RESOURCES

However, the AGV receives its route from the public cloud, which can be connected via the Internet. This is where the planning platform is located, where companies define routes for their entire fleet of AGVs. And that is done centrally for all the locations where automated transport and forklift trucks are deployed.

By exchanging data with the public cloud, companies not only keep track of the position of their vehicles, but can also plan further steps with the data that is generated. For example, in maintenance: It's possible to identify when the vehicle has to go to the charging station more frequently – a strong indication of wear or damage to the battery. Worn parts can be replaced or repaired to prevent breakdowns. In this way, the public cloud and edge nodes become a well-coordinated team.

Providers of AGVs, however, can benefit from thousands and thousands of anonymized pieces of device data – subject to user consent – which they can use for the further development of their products and analyze in the public cloud.

QUALITY CONTROL: LIGHTNING-FAST REACTION WHEN IT COUNTS

Digitalized quality control is another application scenario that many companies are already implementing: Edge nodes monitor the data stream from sensors in production and ensure that predefined threshold values are adhered to. If necessary, the edge system stops production in the same fraction of a second that a fault is detected. It's a mechanism that prevents defective parts from circulating and potentially causing consequential damage. This process must take place in real time – production cannot wait for feedback from a cloud that may come with a few seconds delay.

However, the public cloud is ideal for the analysis of large amounts of data and predictive maintenance. For this reason, the data stream also migrates continuously to the public cloud, where it can be stored inexpensively and analyzed as needed. For example, for machine learning processes: The combination of edge computing and the public cloud in the edge cloud offers the ideal conditions for all applications that require fast response coupled with high scalability and a wide range of functions. At a fraction of the cost that would be incurred for on-premises instances. Small edge computing units are already available at 2 percent of the cost of a private cloud instance. And by combining them with the public cloud, they also offer the same benefits.

LEARN FASTER: MACHINE LEARNING WITH EDGE COMPUTING AND IAAS

In addition, edge computing can significantly accelerate artificial intelligence processes. In machine learning, for example, pre-filtered data streams make machine learning much faster. Relevant information is filtered out before it is sent for analysis. For example, edge nodes can be programmed to analyze video data from surveillance cameras while ignoring redundant images and transferring only those that show changes to the cloud. Then in a second step, the almost infinitely scalable cloud resources then take over the actual machine learning process – which, however, is much more efficient because pre-structured and filtered data is already available.

It's a principle that can also be transferred to other areas: "For example, an edge device on a gas turbine calculates the audio signal of a microphone in real time into the frequency spectrum in order to transport only these small amounts of data to the backend," was the description in a recent article from Crisp Research[5] on edge computing.

CLOUD-EDGE COMBINATION: CENTRALIZED CONTROL OF EDGE NODES

Another advantage of edge computing is the ability to centrally distribute new software and tasks to all the affected sites. For example, edge nodes on-site and the public cloud in the data center share the work: Edge nodes collect data from sensors on machines, pre-filter them, sort out unimportant or redundant data and send the relevant data to the public cloud. There, the data from all locations is analyzed centrally. On the basis of the knowledge gained, companies can use the public cloud to adjust the parameters for controlling their production lines anywhere in the world in real time.

Similarly, companies can provide their edge nodes with software updates centrally from the public cloud. Or give them completely new tasks: So, for example, an edge node that analyzes images from surveillance cameras and forwards irregularities to the cloud today could take over the control of AGVs tomorrow.

CONCLUSION: CLEVER COMBINATION OF EDGE AND IAAS BRINGS SPEED

Even today, lean hardware resources on-site help companies to carry out latency-critical processes reliably and quickly. The combination with the public cloud provides companies with many more opportunities for deriving added value from data generated in the value chain. No in-depth expertise is really required: With the Open Telekom Cloud Edge, companies not only receive the necessary hardware resources, but also expert advice from the relevant experts.

"According to a recent forecast by market analyst firm IDC, in three years' time 40 percent of cloud installations will already have edge computing capabilities.[4]"

DELEGATE CUMBERSOME TASKS
MANAGED HYBRID CLOUD SERVICES

Hybrid cloud infrastructures can offer companies the best of both worlds: Low latency, high security and full control thanks to dedicated capacities (private) on the one hand, scalable resources for peak loads, on-demand billing and access to the latest services at all times on the other (public). However, a huge pitfall can already lurk in the construction of hybrid infrastructures. This is because companies are setting a decisive course for the subsequent management effort. Those who design public and private clouds with different hardware and software components, for example, need qualified employees for both instances who are familiar with both the peculiarities of the public cloud instance and the specifications of the on-premises resources.

A hybrid cloud with a uniform hardware and software basis is more efficient. This is because there is only one technology that has to be served; both for on-premises and in the public cloud. One example of such a construct is the Open Telekom Cloud Hybrid Solution: Telekom supplies companies with hardware and software from the private instance to the required extent, takes care of setup and installation, and also handles operation and maintenance throughout the entire lifecycle.

DECISIVE: USABILITY ACROSS CLOUD BOUNDARIES

This results in a considerable comfort factor in the user experience (UX), which plays a central role when it comes to efficiency in dealing with the IT environment. For the operating teams of a hybrid cloud infrastructure, who work across the board by switching back and forth between public and private as a matter of course, a uniform UX can save considerable time. Ideally, a hybrid cloud infrastructure should be operated via a single console with the same Graphical User Interface (GUI), in which on-premises and public clouds seamlessly merge. One console and one world that combines all the advantages of public and private clouds.

ENSURE OPERATIONAL SAFETY AROUND THE CLOCK

But IT management isn't just about orchestrating resources via a console. Among other things, companies must also ensure that their systems meet the highest possible security standards. This is another reason why it can be worthwhile to outsource operations and maintenance completely to a trustworthy service provider. Many companies operate their own servers around the clock, which run even when the IT department has long since finished work. A malfunction at night, which can lead to damage or loss of data, isn't noticed until the next morning. Managed services from a provider such as Telekom, which monitors the operating status and provides 24/7 standard support, enable companies to ensure maximum reliability even for on-premises resources within a hybrid cloud infrastructure.

For example, in the event of a hardware failure: If a company operates an on-premises instance managed by Telekom, a hardware defect is immediately detected by Telekom, regardless of the time of day. In such a case, an exchange is initiated immediately. Malfunctions in the network or other management components are also registered and, if necessary, rectified as quickly as possible via remote access. Even the company's own applications can be managed by the provider around the clock, if necessary, in order to rectify downtimes as quickly as possible.

MANAGE PAAS AND SAAS, TOO

And what about standard software? After all, more and more companies are not only using naked infrastructure from the cloud, but are also leasing applications – whether office packages, development platforms or entire SAP instances. In this area, too, companies can cooperate with providers such as Deutsche Telekom. The advantages are the same: less effort, higher security and more time for the core business.

Deutsche Telekom, for example, offers numerous standard services, such as SAP systems, which it also operates for companies on demand. The Open Telekom Cloud was certified by SAP for operation in the IaaS data centers in Magdeburg and Biere in Saxony-Anhalt. In addition, an SAP operating model is also available for the on-premises instance of the Open Telekom Cloud for companies, which can also be managed by Telekom if required. A construct that Telekom can offer for every standard service from the partner ecosystem. In this way, providers like Telekom are able to free companies almost completely from IT management tasks.

COMPARISON OF ARCHITECTURES
HYBRID CLOUD VARIETIES

Who is allowed to access which data and when? How large must the dimensions of which systems be? And how do IT managers keep an overview? Designing a hybrid cloud infrastructure is a complex challenge. After all, every industry and every company have different requirements and different budgets: Off-the-shelf solutions don't exist here and wouldn't make sense either. What's more, not all companies have in-house specialists who are capable of designing a complex hybrid cloud infrastructure. For this reason, more and more companies are turning to the expertise of providers such as Telekom. Telekom helps companies design, build and operate such IT topologies.

Companies need to balance costs against performance, complexity, security, and latency. The four most common topologies that enterprises can use to design hybrid cloud architectures are: public front, private front, private isolated and clean sheeting.

PUBLIC FRONT: PRIVATE CLOUD IN THE BACKGROUND

The public front topology is suitable for companies that on the one hand want to benefit from the scalability and flexible costs of the public cloud, but on the other hand don't want to give their customers access to their private cloud instance. External requests automatically end up in the public cloud. Requests are only forwarded from there to the private cloud when necessary. Direct access to the private instance is reserved for company employees via a secure connection. And what is the reason for this? The public front topology enables companies to process particularly sensitive data on their own on-premises resources without having to forego the advantages of a scalable public cloud.

They can also use this construct to protect their systems from unexpected traffic peaks. If requests via Internet Protocol (IP) end up in the public cloud, their own network – the private cloud – is spared this traffic. If there are suddenly a lot of requests, these can be easily intercepted in the public cloud using the freely scalable resources available there. As a result, companies only use and pay for exactly as much as they need for their web frontend and backend and at the same time benefit from maximum security and low latency thanks to their own resources in the private cloud. On-premises resources can be reduced to a minimum, keeping overall costs low.

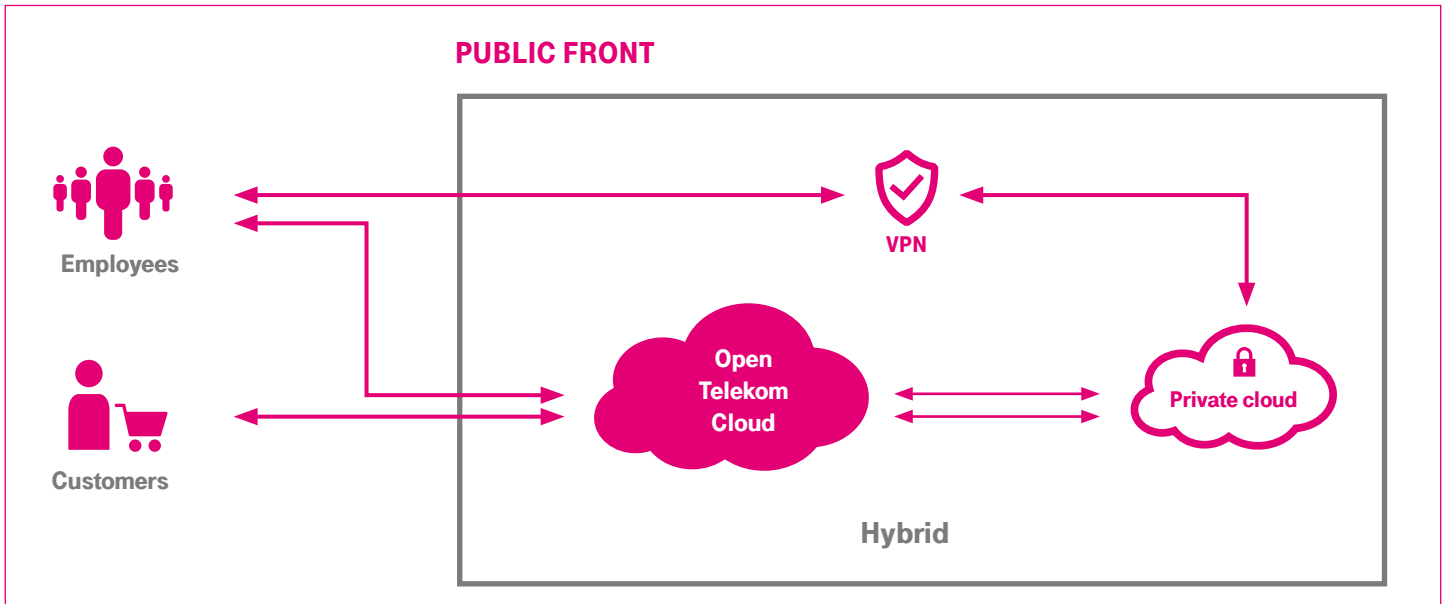

Public front: External requests end up in the public cloud, only company employees have access to the private cloud.

APPLICATION EXAMPLE OF PUBLIC FRONT: BOOKING PORTALS ON THE INTERNET

If you are looking for information about things like train connections, transfer times or the current traffic situation, you can get it from the public cloud, because that's where the web frontends and backends are located. Even in the event of an unexpected number of requests – for example in the event of a rail strike or thunderstorms – the website is able to cope because the resources in the public cloud scale automatically with the demand.

It's only when someone wants to book a ticket that the private cloud comes into play. This is where personal data is stored and processed, as well as sensitive data such as billing information or company-internal data. It would be similar, for example, with a cinema booking portal. Here, too, all requests that generate high traffic can run into the public cloud. For example, the hosting and playback of movie trailers. Only the booking process, during which sensitive personal data and payment information is generated, is handled using the private cloud.

PRIVATE FRONT: KEEP TRACK OF TRAFFIC, HIGH SECURITY

Sometimes, however, exactly the opposite scenario makes sense: If unexpectedly high traffic can be ruled out, private front models can also offer advantages. For example, companies that run all IP requests through their on-premises resources have the best possible overview of where the different requests are coming from – and therefore can control all network access. The advantage: maximum information and control. With this topology, companies can track in real time who retrieves what information and when. And in the event of unwanted requests – such as Distributed Denial of Service (DDoS) attacks – they can block or redirect the corresponding IP addresses.

In addition, companies that use a private front topology meet the highest security and compliance standards. This prevents customer data from entering the public cloud, even temporarily, as is the case with the public front topology. This model could, for example, be of interest to manufacturers of machines that are networked with a platform hosted in the private cloud. Because data from machines that is assigned to specific customers is also legally considered personal data, providers prefer to store it in the private cloud. And they can do so knowing that there will be no issues related to data protection and that the amount of data remains calculable at all times: Anyone who receives machine data can count on a constant flow of information without any sudden traffic peaks.

APPLICATION EXAMPLE OF PRIVATE FRONT MODEL: PREDICTIVE MAINTENANCE

Nevertheless, companies can also benefit from the advantages of the public cloud in the private front model, such as for burst scenarios, analyses or machine learning processes. They can, for example, make customer or machine data anonymous before they transfer it to the public cloud to analyze it or use it to train an AI. This can be data from an ERP system that is used to forecast the purchasing behavior of customers, or data from machines used by customers for predictive maintenance processes.

The disadvantage of the private front topology is the slightly higher cost. That's because, in order to guarantee the availability of their own services at all times, companies need a level of resources in the private cloud instance that is at least equal to the maximum expected use – including external requests. This means that the on-premises share must inevitably be greater than the public front topology, which is reflected in correspondingly higher costs for the non-scalable resources.

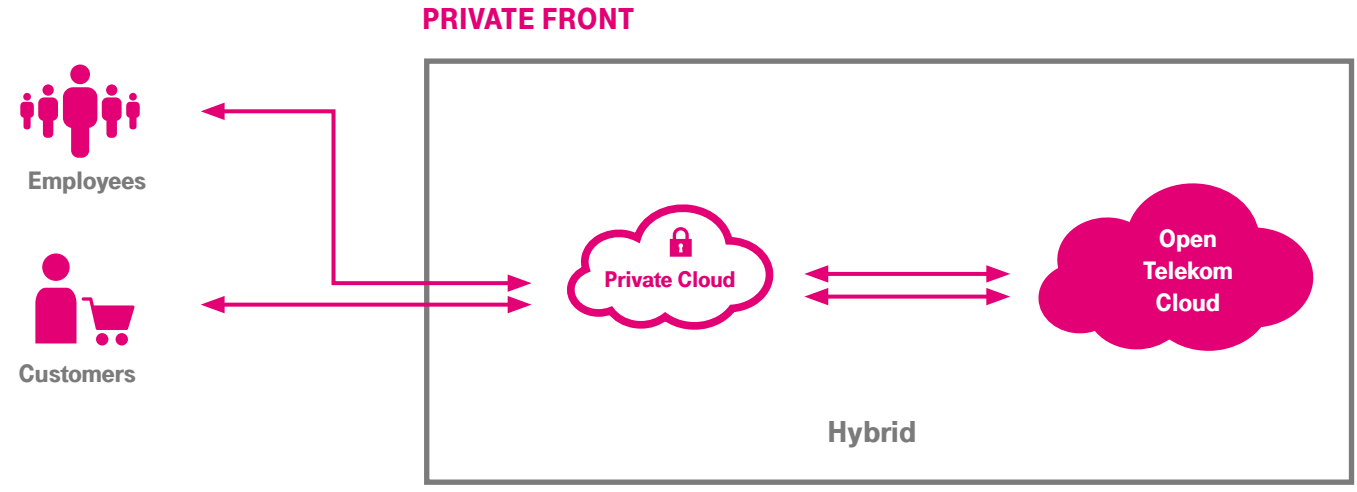

Private front topology: Everything runs on-premises via the private instance.

PRIVATE ISOLATED, HYBRID DISCONNECTED: CONNECTION UNWANTED

However, although the interaction between private and public clouds can work very well, certain application scenarios require a strict separation between the instances. In some industries, for example, there can or must be no connection between on-premises and public cloud instances. Some companies are confronted with this problem because most providers only offer a hybrid cloud solution with an existing connection to the public cloud instance. In Germany, only the Open Telekom Cloud Hybrid Solution currently offers a completely separate operation of public and private cloud resources on the basis of an identical hardware and software architecture.

APPLICATION EXAMPLE OF PRIVATE ISOLATED: FOR PROTOTYPES AND WIND TUNNEL TESTS

By operating the instances in isolation, companies can handle customer requests via the public cloud. Automotive companies use the public cloud for less sensitive processes such as hosting their website, including a car configurator. However, confidential company data such as the construction drawings of prototypes, crash test or wind tunnel simulations never leave the company premises for security reasons and are carried out in the private cloud – without any connection to the outside world.

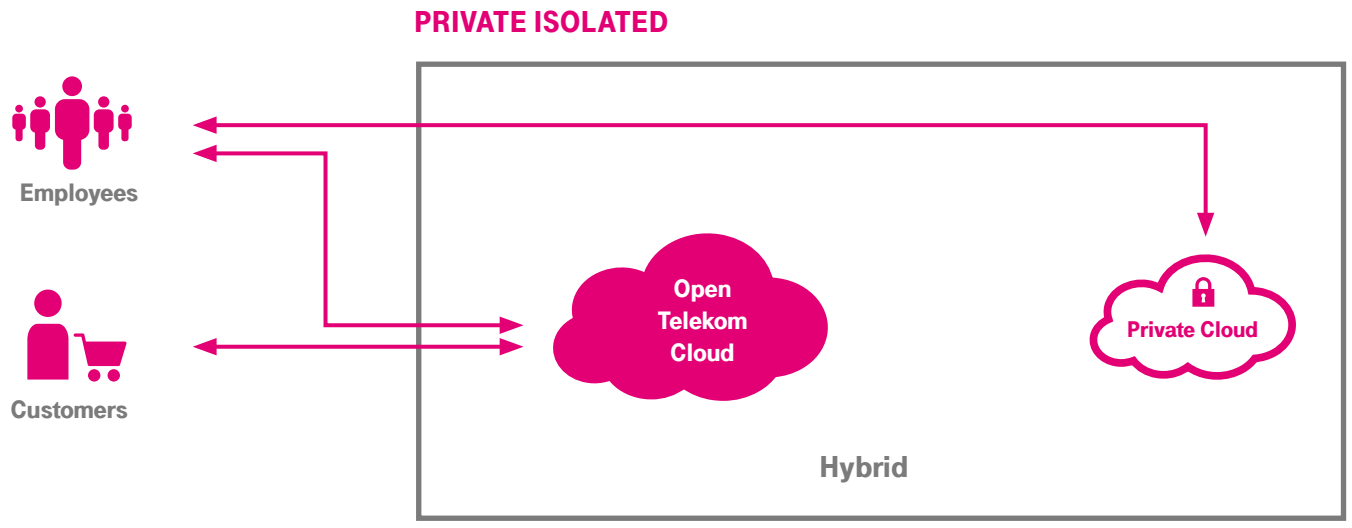

Private isolated: No connection between private and public cloud is mandatory in some industries.

CLEAN SHEETING: COMPLEX OPTIMUM

The most complex topology is the combination of all the variants: With so-called clean sheeting, companies create a complex set of rules that shows which user may access which resources at which time. Companies thus combine all the advantages of the different operating concepts. Employees no longer have to worry about where the required resources come from and can concentrate on their core business – both in the company divisions and in IT. Each request is served by an automated broker from the public or private cloud, according to specific parameters, and the most cost-effective operating mode is chosen – taking the policies into account. This saves companies twofold: once through automation, and once again through the most cost-effective operating mode.

Customers are also guided by an automated logic into the private or public cloud, depending on their request. They are only allowed direct access to the private cloud under certain conditions, but companies are still well prepared for sudden traffic peaks thanks to the public cloud. Sensitive data remains in the private cloud.

The difficulty with this topology is the effort involved and the challenge of keeping an overview. In clean sheeting, every mistake can have serious consequences. Designing such an architecture requires in-depth expertise. That's another reason why companies should consider relying on an experienced partner like Deutsche Telekom, which not only takes care of the design, but also subsequently implements the operation and maintenance if necessary, so that all the company's employees can really focus on the core business.

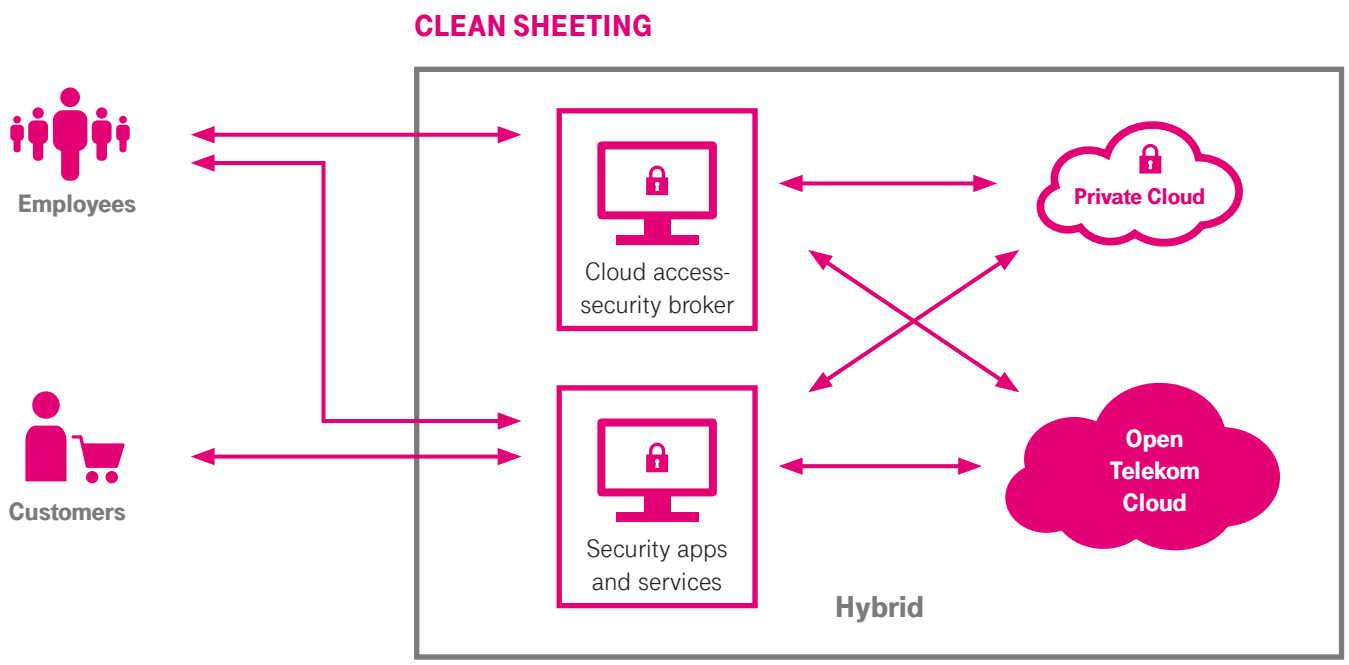

Everyone with everyone: Clean sheeting has many advantages but requires in-depth expertise.

SOURCE LIST

1)
Trend study: Hybrid Cloud in Germany - Needs and requirements in SMEs for the combined use of IT technologies operated in the cloud and on-premises. February 2018. Telekom/CXP Group (BARC – le CXP – PAC)

2)
Data Age 2025: The Digitization of the World. Seagate/IDC Research

3)
OpenStack 2018: International study project in six European core regions. IDG Communications Media AG

4)
Trend forecast "IDC FutureScape Predictions 2019". IDC Research

5)
https://www.crisp-research.com/iot-edge-von-gateway-bis-machine-learning/ Retrieved on 07.01.2019

CONTACT:

open-telekom-cloud.com/en/contact

PUBLISHER:

T-Systems International GmbH
Hahnstraße 43d
60528 Frankfurt am Main
Germany

Date: October 2019

LIFE IS FOR SHARING.