

White paper

On the Fast Track to Your Ideal Virtualization

Today's IT infrastructures must be ultra-agile to meet fast-changing business challenges and opportunities. What's needed is a new and sustainable approach, with efficient and flexible infrastructures that span across servers, networks and storage environments. The answer is virtualization. Experience the various flavors of virtualization and take advantage of FUJITSU Integrated System PRIMEFLEX to quickly transform your IT infrastructure without risks and in a cost-effective way.



Content	
Virtualization started long time ago	2
The first exercise was about servers	3
Bringing the benefits of server virtualization to desktops	3
Storage virtualization – separating data management from the physical storage zoo	4
Network virtualization	4
Software-Defined Data Center (SDDC)	4
An efficient way to virtualize: Hyper-converged infrastructure	4
Classical or hyper-converged? – Use cases matter	5
FUJITSU Integrated System PRIMEFLEX – Your fast track to data center infrastructures	6
PRIMEFLEX meets virtualization	6
FUJITSU Integrated System PRIMEFLEX for VMware vSphere	6
FUJITSU Integrated System PRIMEFLEX for VMware vSAN	7
FUJITSU Integrated System PRIMEFLEX for VMware Cloud Foundation	7
FUJITSU Integrated System PRIMEFLEX for Microsoft Storage Spaces Direct	7
Positioning	7
Summary	8

Virtualization started long time ago

As long as there have been data centers, their complexity often caused sleepless nights for the data center managers. The root cause of this high complexity was the tight coupling of the layers in the data center infrastructure solution stack, such as hardware, operating system and applications. The consequence of this tight coupling was a strong dependency between the layers. Any change in any of these layers impacts its adjacent layers, too. Changing for instance the operating system version, you have to verify if the new version can still run on the existing server hardware, while verifying at the same time whether the applications that were running on the previous operating systems version are still supported by the new one. And unfortunately, changes in any of these layers happen pretty often.

Of course, if the tight coupling is the root of all evil, the solution is quite obvious: You just break the tight coupling between the layers to make the layers and their lifecycles independent from each other. This

is exactly what virtualization stands for. By introducing a so-called virtualization layer between two layers in the infrastructure solution stack, you hide the complexity of the layer underneath. At the same time, this virtualization layer mediates the access to the layer below, supports the creation of virtual resources, it maps virtual resources to physical ones, while providing a strict isolation between the virtual resources.

Virtualization is basically applicable to all layers of the solution stack, be it hardware (server, storage, and network), operating system or applications.

It is well known that virtualization is nothing new. Virtualization had its origin long time ago in the mainframe era, in the so-called 1st platform. Later it was applied to the 2nd platform (client / server architectures), and of course today nobody can imagine the 3rd platform which is about Mobile, Social, Big Data and Cloud and the likes, without virtualization.



The first exercise was about servers

The first exercise regarding virtualization was about servers. The idea of server virtualization is to abstract the operating system from the server hardware, which in other words makes the operating system independent from the physical server hardware. The virtualization layer between the server hardware and the operating system is called hypervisor, also known as virtual machine manager. The virtual resources above the hypervisor are denoted as virtual machines. The hypervisor ensures that compute, storage and I/O resources are shared between all virtual machines.

The benefits of server virtualization are versatile. It allows you to consolidate your applications on fewer physical servers, which helps reduce data center footprint, energy consumption and cooling requirements. Utilization of available resources is dramatically improved. New applications can be provisioned extremely fast, as there is no need to configure, to select, to order and install a physical server every time a new app is deployed. Moreover, server virtualization gives you more flexibility with regard to how you use your resources. This enables you to rapidly react to dynamically changing demands. And you can easily scale your infrastructure on demand to get the performance you need. The built-in high application availability reduces downtime and contributes to business continuity in a cost-effective and affordable manner. Due to the independence of the guest operating systems from the server hardware, the replacement of hardware is an easy thing to do. The independence even enables you to run an old operating system on latest server hardware, thus extending the lifetime of your legacy applications. Running your applications in dedicated virtual machines helps avoid any interferences and conflicts between the applications. Many of the advantages described are essential contributions to reducing complexity, administration efforts, as well as operational but also capital expenditure. At the same time, server virtualization is an essential step to make your IT infrastructure more agile.

Through workload isolation, you run applications previously running on the same physical server in dedicated virtual machines, thus evading any interference between the applications. Workload aggregation, also known as server consolidation, stands for running applications previously running on dedicated physical servers in virtual machines, thus tremendously reducing the number of physical servers needed. Workload migration denotes the capability to run virtual machines on any physical server of your infrastructure, while always getting access to their data on the storage system involved. This helps balancing workloads to get a better overall performance behavior, and react to server failure to avoid downtime.

Bringing the benefits of server virtualization to desktops

It is a matter of fact that organizations using server virtualization could capitalize on it and were wondering if anything similar could be applied to desktops. This was the beginning of the Virtual Desktop Infrastructure, also known as VDI. VDI allows desktop operating systems including desktop applications to be run as virtual machines (virtual desktops) on a physical server infrastructure in the data center. Instead of a fully equipped PC, a thin client or any other device is sufficient for the access. This means that the network becomes much more important than ever before.

The centralization of desktops enables a flexible access to your desktop environment from any device. This even supports trends such as BYOD (Bring Your Own Device). Desktop management and maintenance become much easier, because nearly everything is in the data center. Software can easily be deployed and updated, and patches become effective without touching numerous end user devices and disrupting the users. Onsite visits for end user support are a thing of the past. By running desktop applications in isolation (dedicated virtual desktops), compatibility conflicts can be fully excluded. The level of application and desktop availability is significantly increased; even disaster recovery concepts known from the server world can be applied to end user workplaces. Furthermore, the flexible workplace infrastructure helps react rapidly on new or changing requirements, e.g. in conjunction with mergers and acquisitions which nowadays happen on a regular basis. The fact that all data is hosted centrally minimizes the risk of data theft. Data backup no longer depends on whether the device is turned on or whether it is connected, thus minimizing the risk of data loss and improving security. All this helps meet compliance demands much better. Although additional infrastructure is required in the data center, the overall energy consumption can be reduced, if thinner low-current device hardware is deployed. In the event an access device breaks, it can simply be replaced by a new one.

Virtual desktops can be persistently assigned to end users. However, considering that not all users require their virtual desktops at the same time, it can make sense to create a pool of virtual desktops sized for the maximum number of simultaneously active users. This saves data center resources and further reduces cost. End users are authenticated and dynamically connected to their virtual desktops by a connection broker. If at logon time the personal environment is not available in a virtual machine yet, an available virtual machine will be selected from a pool and personalized.

All these advantages lead to higher efficiency and correspondingly to cost reduction. In other words: desktop virtualization helps overcome the limitations of traditional desktop computing.

Storage virtualization – separating data management from the physical storage zoo

In the same way as with server virtualization, a server operating system is decoupled from server hardware, data services, such as data deduplication, data compression, data replication, snapshots and the like, no longer run in intelligent proprietary storage controllers. They are abstracted from the storage hardware and run on standard servers instead. Thus, the data services become agnostic to the storage hardware. This is exactly what storage virtualization is about. Storage virtualization enables pooling of storage capacities across multiple systems.

Similar to server and desktop virtualization, attractive advantages can be derived from storage virtualization. Complexity is reduced; no more need to deal with storage hardware-specific details; no disparate storage management for different data types; and overall storage deployment and management are significantly simplified. This in turn reduces operational expenditures. Other strong points are flexibility and scalability; it is so easy to add data volumes without disrupting running applications.

Network virtualization

Server and storage virtualization certainly contribute to increasing speed and flexibility in your data center. However, an important aspect is still missing: The network which connects all of the servers and storage systems in your data center. It has to cope with accesses to and from the outside world, from both mobile and stationary devices, while ensuring the usual top quality of the user experience, regardless of where users are located. Another major challenge for the network is handling the huge workload stemming from the Internet of Things (IoT). This makes it obvious that we need the same flexibility and level of automation for networks as for server and storage operation.

In traditional networks, the control logic and the data flow are implemented in the same network device, typically something proprietary from any network component vendor. Therefore network configuration has to happen manually at the device level, which is not only complex, costly and cumbersome, but also inflexible. This makes it impossible to quickly adjust the network configuration to new requirements.

Network virtualization decouples the control logic from the data flow. While data is forwarded locally through the network device (which need no more be as intelligent as in traditional networks), management happens centrally. It is no longer up to network device to find a route from the source to the destination of a data flow; an optimal routing can rather be controlled by intelligent, central management software, even in an automated way. This makes network operation really easy, increases operational efficiency, enables a rapid deployment of new services, and reduces operational cost up to 80%. If new service requirements are to be met, a reconfiguration may be done by just a click of the mouse. Likewise, it is easy to scale your network infrastructure on demand by just adding or moving the network devices; everything else is defined by software. By maximizing the network value, you make your data center more responsive to business needs. And not to forget: Network virtualization makes it much easier to increase the level of security for all your data traffic.

Software-Defined Data Center (SDDC)

With all flavors of virtualization, the key to success is based on the fact that resources can be defined by software. With server virtualization it is the compute resources (Software-Defined Compute, SDC), with storage virtualization it is the storage resources (Software-Defined Storage, SDS) and with network virtualization it is the network resources (Software-Defined Networking, SDN) which can be managed in a highly efficient manner. Using all these options in parallel leads to the Software-Defined Data Center (SDDC).

As in a SDDC everything is virtualized – compute, storage and network resources, any standard server hardware, any storage and any network components can be used. These components are interchangeable at any time. A unified management for all types of resources, even across various sites, tremendously simplifies data center management, requiring less specialist skills and less training.

An orchestration platform controls the quick provisioning and configuration of all resources, manages failure situations and automates the fulfillment of pre-defined service levels by for instance adapting your infrastructure in a way that your workloads always get the resources they require. This helps avoid human errors, gives you the agility and speed you need to quickly respond to changing business demands, and gives you much more time for innovation.

And not to forget the capability of scaling, which is essential for digitalization, especially when considering the ever increasing data volumes which have to be collected, transported through the network, processed and stored.

All told, a Software-Defined Data Center (SDDC) provides all the advantages of Software-Defined Compute (SDC), Software-Defined Storage (SDS) and Software-Defined Networking (SDN), but even more on top. It ensures business continuity, improves service quality, increases agility and reduces cost. It represents a radical paradigm shift for IT, turning your IT organization from a reactive service provider to a proactive change agent.

A Software-Defined Data Center can be easily expanded to a cloud infrastructure by adding a self-service portal, as well as metering and charging functions to its management instance. This will then enable the seamless orchestration of on premise resources but also the resources of external cloud providers enabling Hybrid IT Operation.

An efficient way to virtualize: Hyper-converged infrastructure

Classic data center infrastructures consisted of servers, storage systems, network components and software. By virtualizing these components you reduce complexity, but you still have to deal with all these components. With hyper-converged infrastructures, this looks different. They tightly integrate all compute and storage resources in a commodity server node, making a dedicated physical Storage Area Network (SAN) with its management superfluous. Instead, storage spreads across the local disks of the server nodes. As there is no external storage included, data center footprint will often be reduced just as energy consumption and cooling requirements. Having compute and storage resources integrated in a single box makes deployment even easier and faster.



The built-in data services, such as data replication, snapshots, deduplication and storage tiering turn hyper-converged systems into a software-defined storage platform. The unified management for both compute and storage resources brings simplification to a new level by reducing administration efforts and skills demands.

While classic infrastructures scale on a component level, hyper-converged systems enable scalability on a system level. Compute performance and storage capacity can be scaled by just adding or removing servers. Thus, hyper-converged infrastructures can be easily aligned to growing business demands, while business continuity is always ensured. And of course, all benefits aforementioned have often a positive impact on capital and operational expenditure.

Classical or hyper-converged? – Use cases matter

With all of the benefits previously mentioned, you might come to the conclusion that a hyper-converged infrastructure will be the only way to go in the future. But let us just consider a number of aspects which might help you in making the right choice.

If your workloads scale horizontally, hyper-converged will be a perfect fit, especially if compute and storage resources need to scale in tandem. This applies to workloads which require a fixed amount of CPU performance, main memory, disk space and IOPS. Typical examples are Hosted Virtual Desktops and Hosted Shared Desktops. If your workloads scale vertically, or require granular expansion at the component level, hyper-converged might be less appropriate.

Virtualization is a prerequisite for hyper-converged infrastructures. In other words, these infrastructures cannot be used for workloads which run on bare metal only – a virtual environment would be ineffective and slow them down. And since most hyper-converged implementations are based on a single hypervisor, they will not be a good fit in cases where mixed operation of multiple hypervisors is needed to run different workloads.

Hyper-converged concepts have become attractive options for remote offices and branch offices. As no external storage infrastructure needs to be maintained, frequent costly onsite visits can be avoided. There are customer cases where travel time could be reduced by 99 percent just by replacing a physical SAN with a hyper-converged infrastructure.

If your workloads benefit from data services coming from hyper-converged scenarios, you may use these services without any additional investment. If you don't need them, you will pay indirectly for things you don't use. Another aspect worth considering is expected growth. The more frequently you have to expand your infrastructure, the more you will benefit from the ease of scalability available through hyper-converged infrastructures.

The unified management of compute and storage resources reduces operational complexity, administration work and costs. But bear in mind that adopting this new approach will change existing staff roles and require other organizational changes. You could face some resistance from your IT staff, especially in the storage area. Will you be able to counter this resistance? This aspect is also closely related to your storage strategy. If you intend to utilize existing storage, hyper-converged solutions will hardly fit into your strategy. However, if you intend to replace your existing storage sooner or later, opting in favor of a hyper-converged solution might be a good start.

The storage capacity of a hyper-converged infrastructure is limited by the number of server nodes. If you have to cope with data volumes larger than the maximum storage capacity of your server cluster, hyper-convergence will not be a feasible option. Even though the hyper-converged approach promises linear scalability, predictable network performance with very large deployments is still an open question due to a lack of experience in this area.

In addition to the technical characteristics of relevance to workloads, software licensing aspects should also be taken into account. For instance, a database application may be a perfect fit for a hyper-converged scenario, but if you have to pay license fees per CPU socket or even per CPU core, the hyper-converged approach will probably not be feasible from a business standpoint.

Ultimately it is all about cost. As mentioned before, operational expenditures always tend to be much lower for hyper-converged infrastructures when compared to classical converged concepts. But when it comes to capital expenditures, it is hard to make a general statement. Typically, from a hardware cost perspective, hyper-converged is more attractive than classic; from a software cost perspective, it is just the reverse. Hyper-convergence requires a minimum number of server nodes, plus special certified hardware components, and license fees will be charged for the virtualization software. You will find lots of examples with cost advantages on either side. Make a simple cost comparison for your specific project to find out which option is most cost-effective for you.

It is the use case which matters when comparing classical converged and hyper-converged scenarios. There are good reasons to look at both architectural approaches. You should make final decisions based specifically on each use case. Whenever their benefits outweigh the drawbacks, you should go for hyper-converged systems.

FUJITSU Integrated System PRIMEFLEX – Your fast track to data center infrastructures

No matter which flavors of virtualization are applied, and no matter if you are going for classical or hyper-converged, your data center infrastructure needs to be built up. But building a data center infrastructure can be extremely complex. Servers, storage, network components and software need to be selected from a myriad of options, procured and integrated. As the compatibility of the components is not guaranteed at all, extensive testing is a must. This in turn requires a deep knowledge of all components involved and an understanding of their interdependencies on each other. Consequently, a do-it-yourself approach will be time-consuming and expensive, while presenting businesses with multiple risks that the final infrastructure won't work as desired.

There is a better and more cost-effective way to go: FUJITSU Integrated System PRIMEFLEX, a pre-defined, pre-integrated and pre-tested combination of data center components, such as servers, storage, network connectivity and software. As all typical activities, such as infrastructure design, component integration and testing have been completed before starting the project, onsite activities only include the deployment of the Integrated System and the integration into the production environment.

PRIMEFLEX systems are built from best-in-class components, either Fujitsu's own technologies, such as our FUJITSU Server PRIMERGY and FUJITSU Storage ETERNUS, or from leading technology partners. Based on best practices and real-life project experience, PRIMEFLEX systems are designed in a way that their components will work optimally together. What is more, they are proven in Fujitsu's cloud operation.

PRIMEFLEX systems are either pre-installed in the factory, and arrive ready-to-run at the customer's site; or they are delivered as reference architectures giving you the flexibility to adapt them to your specific requirements. On demand the adjusted configuration can be pre-installed and delivered ready-to-run, combining the advantages of reference architectures with those of ready-to-run systems. For all PRIMEFLEX reference architectures, installation and configuration guidelines are available as a standard.

The PRIMEFLEX family includes both, systems based on a classical and systems based on a hyper-converged architecture. And PRIMEFLEX is supplemented by services throughout all lifecycle phases, either delivered by Fujitsu or its local partners.

The benefits resulting from the PRIMEFLEX Integrated Systems approach are manifold. Due to reduced complexity, introducing a new infrastructure in your data center becomes much simpler. You will experience less trouble through trial-and-error testing, because the compatibility of all components is always guaranteed. At the same time, risk is minimized and the skill sets required in your IT department will be less demanding. Apart from this, less time is needed for planning, and deployment is tremendously accelerated which shortens time to production and time to value. Due to the optimized design, resource utilization is optimized. This can have a positive impact on data center space, cabling, energy consumption and cooling efforts. Moreover, an Integrated System represents a perfect foundation for efficient operations and reduced maintenance efforts. All these aspects help reduce cost, both CAPEX and OPEX. Finally, we should not ignore the fact that all these benefits enable IT organizations to focus on the really important aspects of the business. Moving away from a build and maintain focus means improved responsiveness to new business requirements, or even driving business to a new level.

PRIMEFLEX meets virtualization

With our PRIMEFLEX family, we address highly relevant use cases. One of the focus areas is virtualization. Powered by FUJITSU Server PRIMERGY, PRIMEFLEX benefits from all its proven values. PRIMERGY is the best performing server platform for running VMware virtual machines. Their long and proven track record is reflected in top rankings across all VMware VMark benchmark categories. Its patented Cool-Safe® technology drives world records for energy efficiency. Moreover, PRIMERGY is shown as a leading platform in global server reliability reports.

Subsequently we are showing a brief overview of the individual PRIMEFLEX offerings for virtualization and their positioning.

FUJITSU Integrated System PRIMEFLEX for VMware vSphere

PRIMEFLEX for VMware vSphere is based on the classical infrastructure approach with external storage from FUJITSU (either ETERNUS AF or ETERNUS DX). Network switches, cabling and the rack infrastructure are also included. The built-in Infrastructure Manager ISM from Fujitsu supports easy converged lifecycle management for all components. Customers can choose from standard configurations in various sizes. These configurations can easily be customized. PRIMEFLEX for VMware vSphere is delivered as a ready-to-run solution.

→ For more information, [click here](#).

PRIMEFLEX, jointly developed and marketed by Fujitsu and NetApp, is a converged infrastructure alternative that includes NetApp storage. Various configurations based on PRIMERGY CX servers, NetApp storage (NetApp FAS in the entry configuration and NetApp All-Flash FAS in the others), VDX switches from Extreme Networks and the Fujitsu Infrastructure Manager ISM are fully integrated at the factory and shipped to customers as ready-to-run solutions. Various expansion packages are available to ensure easy scalability.

→ For more information, [click here](#).

FUJITSU Integrated System PRIMEFLEX for VMware vSAN

PRIMEFLEX for VMware vSAN is a hyper-converged system based on VMware vSphere and VMware vSAN, supporting up to 64 server nodes. Various configurations for specific use cases are in place, such as All-Flash configurations for write-intensive workloads demanding for low latency (e.g. virtual desktop infrastructure with either linked clones or full clones), hybrid configurations with hard disks and Solid State Disks for mixed workloads, and special high density configurations for use cases with data center footprint being the critical parameter. The reference architecture approach allows flexible adjustments of the pre-defined configurations with regard to processor type, main memory size and storage capacity. PRIMEFLEX for VMware vSAN comes with VMware vSphere and vSAN pre-installed.

PRIMEFLEX for VMware vSAN are also a sound foundation for private cloud infrastructures.

→ For more information, [click here](#).

FUJITSU Integrated System PRIMEFLEX for VMware Cloud Foundation

PRIMEFLEX for VMware Cloud Foundation represents a truly Software-Defined Data Center with everything virtualized based on a hyper-converged architecture. It is based on VMware vSphere, vSAN and NSX for virtualizing compute, storage and network resources. The additional SDDC Manager acts as an automation engine for provisioning, monitoring and lifecycle management. VMware vRealize is an optional add-on that turns PRIMEFLEX for VMware Cloud Foundation into a full private cloud infrastructure. A further option is VMware Horizon which makes the platform usable as a virtual desktop infrastructure.

As the Integrated System is delivered ready-to-run, it is also ready-to-operate within a couple of hours after arrival on-premises. It is available in various configurations from 4 to 240 servers.

→ For more information, [click here](#).

FUJITSU Integrated System PRIMEFLEX for Microsoft Storage Spaces Direct

PRIMEFLEX for Microsoft Storage Spaces Direct is a hyper-converged system based on software-defined storage technology (Storage Spaces Direct) integrated in the Windows Server 2016 Datacenter Edition. The system supports up to 16 server nodes. Various configurations for a broad range of use cases are in place covering mixed workloads as well as workloads requiring extreme I/O performance. Configuration options include hard disks, Solid State Disks and high-speed NVMe disks that allow for setting up a 2-tier and 3-tier storage infrastructure. The reference architecture approach allows flexible adjustments of the pre-defined configurations.

→ For more information, [click here](#).

Positioning

All PRIMEFLEX systems aforementioned can be used for virtualization scenarios. To identify the ideal one for you, multiple aspects need to be considered. Here are the most important questions to be answered:

- On which operating systems are your applications running?
- Which hypervisors support these operating systems?
- What's your preference, if several hypervisors meet the needs?
- If you are going for VDI, do you have any preference in terms of VDI software?
- Is a classic or hyper-converged architecture more appropriate?
- How many physical servers need to be virtualized?
- What are the growth resp. scalability demands?
- Is a disaster recovery option needed, and if so, which one?
- What are the timely expectations on getting the solution up-and running?

The following table will guide you to the virtualization solution which is ideal for your needs. And the good message is that with PRIMEFLEX you will always have the right solution whatever your requirement will be.

	PRIMEFLEX for VMware vSAN	PRIMEFLEX for VMware Cloud Foundation	PRIMEFLEX for Microsoft Storage Spaces Direct	PRIMEFLEX for VMware vSphere	NFLEX
Architecture	HCI	HCI	HCI	CI	CI
Server	PRIMERGY RX/CX	PRIMERGY RX	PRIMERGY RX	PRIMERGY RX	PRIMERGY CX
Storage	Local disks	Local disks	Local disks	ETERNUS DX/AF	NetApp FAS/AFF
Networking	Any	Extreme Networks, Cisco (optional)	Any	Extreme Networks, Broadcom (optional)	Extreme Networks
Hypervisor	VMware	VMware	Microsoft	VMware	VMware
SDN	No	Yes	Yes	No	No
Infrastructure Management	ISM (optional)	ISM	ISM (optional)	ISM	ISM
Scalability	4-64 servers	4-240 servers	2-16 servers	2-64 servers	2-64 servers

Summary

Virtualization is an important step to reduce complexity in the data center. But it is no longer sufficient to apply it to servers only. Virtualization can provide the same benefits to desktops, to storage, and to networks. In the era of digitalization where you have to deal with huge amounts of data and quickly react to ever changing demands, virtualization across the board with a high level of automation is needed. Therefore the Software-Defined Data Center resounds throughout the land.

Hyper-converged infrastructures represent a highly efficient way to go for virtualization and provide many additional benefits compared with virtualization in a classical architecture. But they are not the Holy Grail for every situation, because use cases matter.

FUJITSU Integrated System PRIMEFLEX helps you reduce complexity, time, risk and cost when building your virtual data center infrastructure. PRIMEFLEX covers classical and hyper-converged architectures, virtualization for servers, desktops, storage, networks and entire data centers, virtualization technology from VMware and Microsoft, and much more. We want to give you choice and the option to jointly identify what the ideal virtualization solution is for you – A proof-point of Fujitsu's Business-Centric Data Center approach.



Published by
Fujitsu Technology Solutions GmbH
Mies-van-der-Rohe-Strasse 8
D-80807 Munich

www.fujitsu.com/primeflex

2018-01-24 WW EN

© 2019 Fujitsu, the Fujitsu logo, and other Fujitsu trademarks are trademarks or registered trademarks of Fujitsu Limited in Japan and other countries. PRIMEFLEX is a registered trademark in Europe and other countries. Other company, product and service names may be trademarks or registered trademarks of their respective owners. Technical data subject to modification and delivery subject to availability. Any liability that the data and illustrations are complete, actual or correct is excluded. Designations may be trademarks and/or copyrights of the respective manufacturer, the use of which by third parties for their own purposes may infringe the rights of such owner.