

Redundancy concepts for hierarchical switch networks

The issue of high availability is one of the most important aspects when planning for reliable switch networking. Failures as a result of misconfiguration often lead to entire communications infrastructures going down. The consequences include immense follow-up costs and production downtime. With good planning, the redundant connection of the switches across the entire network minimizes those risks of failure and increases the availability of networks.

This techpaper informs you about the most important protocols for redundancy in networks and gives you examples of how a highly available three-tier or two-tier network can appear.

This paper is part of the **series "switching solutions"**.

Click on the icons to find out more about the information available from LANCOM:



Design guide
Redundancy
concepts for switch
networks

The three redundancy concepts VPC, stacking, and STP

By connecting a switch to two different switches in the aggregation/distribution layer or core layer above it, the use of Link Aggregation Groups (LAG) results in extremely high availability (HA) and practically uninterrupted network operations. An important factor here is the use of loop prevention mechanisms. Various redundancy solutions are available for networking two switches, including the Spanning Tree Protocol (STP), which is less effective, and better options such as the Virtual Port Channel (VPC), or stacking.

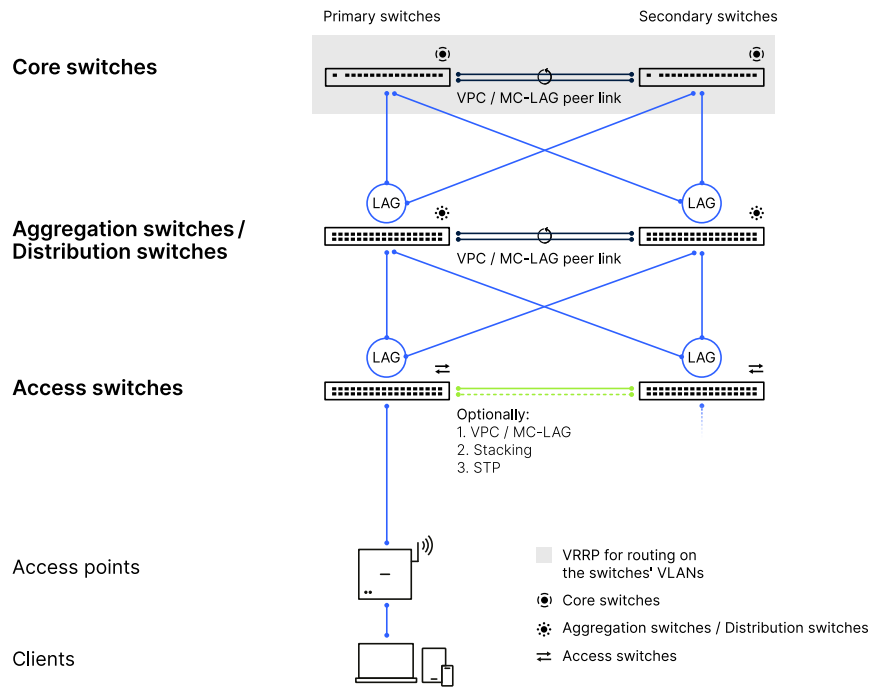



Figure 1:
Three-tier network
with redundancies

The differences between the three protocols VPC, stacking, and STP include the complexity of the configuration, the downtime when restarting the switches, and the cost of the necessary switches.

	VPC	Stacking	STP
Complexity & configuration time	Medium	Low (almost plug & play)	High
Hardware requirements → Costs	High	Medium	Low
Redundancy for high availability	✓	✓	✓
Bandwidth increase & load balancing (active/active)	✓	✓	—
100% network uptime	✓	— (Maintenance window required)	✓
In-Service Software Upgrade (ISSU)	✓	—	—

The three protocols are presented in detail below.



Virtual Port Channel (VPC)

VPC belongs to the Multi-chassis Etherchannel [MCEC] family and is therefore also known as **MC-LAG (Multi-Chassis Link Aggregation Group)**. Due to the high hardware requirements, it is the most cost-intensive of the three redundancy solutions and is therefore usually used in large network infrastructures. To improve failure tolerance through redundancy, this virtualization technology makes two interconnected switches appear as one virtual link. VPC has the following properties:

- **Redundancy and load balancing:** Using their peer link, switches in the virtual VPC group constantly exchange important information about the network, including MAC tables. Each peer switch processes half of the data volume from the access layer (active/active technology). In contrast to stacking, they remain independent instances and it is only the connected ports that virtualize the reciprocal redundancy. However, the administrator needs to configure the VPC peers identically and stacking is not supported on the VPC switches at the same time.
- **100% uptime through rapid convergence:** In the event of a device failure or a change to the network, VPC quickly recalculates the network paths. This eliminates a single point of failure, resulting in faster service recovery. The other device in the VPC cluster handles all of the traffic and keeps the network active. This is irrespective of whether the device failure was caused by a defect or a deliberate shut down, such as during a firmware update (In-Service Software Upgrade, ISSU). This achieves 100% uptime of the network from the core to the end devices.
- **Independent management:** From the perspective of a third device, the peer link makes the switches appear as a single logical-link access point or layer-2 node. The third device can be a switch, server, or other underlying access-layer network device that supports link aggregation. As mentioned above, the peer switches remain independently manageable devices that can be rebooted or updated individually. However, VPC is always controlled by the primary switch. VPC implementations from other manufacturers are usually not compatible with each other.
- **Increased bandwidth:** Bundling the peer link (active/active) increases the bandwidth and throughput capacity between the devices.
- **Simpler network topology:** Because VPC enables LAG between network layers, it reduces the need for STP, which is used in traditional L2 networks to avoid loops.
- **Support for non-VPC enabled devices:** VPC enables end devices or network components that are not VPC-capable to connect to a VPC environment. For this

purpose, these end devices must have a LAG and therefore usually two Ethernet ports (exception: static LAG via one cable).

- **High-performance switch hardware:** VPC places high demands on the switch hardware, which must support the VPC protocol. This can limit the choice of devices, especially at the access layer, and can be costly.




Stacking

A stack is a group of switches that physically behave as a single device. All devices in the stack must have the same stacking interfaces (ports) and be equipped with an identical firmware version. Similar to a chassis or blade system, the stacking ports handle all data traffic in hardware with protocols optimized for this purpose. The stacking technology can be summarized as follows:

- **Almost plug-&-play configuration**
- **Layer-2 simplification:** Stacking can be imagined as a backplane of the individual switches connected via cables that is not recognized as a connection by the configured layer-2 protocols. This allows network traffic to be transmitted over multiple connections simultaneously, so maximizing throughput.
- **No layer-3 routing required:** The intelligent distribution of the data stream within the stack does not require layer-3 routing because the internal stacking protocols handle the connections as described above.
- **Fast failover and almost uninterrupted forwarding:** Thanks to fast detection and link recovery technologies, stack connections are transferred to other switches in the event of failure by means of "hitless failover", i.e. without data loss.
- **No in-service software upgrade:** A disadvantage with stacking is that stacked switches have to go offline during a firmware update, i.e. 100% uptime is not guaranteed during software updates or reboots. Nevertheless, this option can be considered as an alternative to the VPC when maintenance windows are used. During operation, active/active operation achieves maximum data throughput between the core and end-device layers.

STP



Spanning-tree protocol (STP)

The technical differences between the current spanning-tree standards **MSTP** (Multi-STP, IEEE 802.1s) and **RSTP** (RapidSTP, IEEE 802.1w) are not discussed here. Instead we make reference to the relevant literature. While VPC and stacking focus on physical redundancy and load balancing, STP provides a logical solution to avoid network failures due to loops and to ensure fast recovery.

Of the three protocols presented here, STP has the most laborious configuration. Although STP can achieve zero downtime in active/passive mode between the access-switch layer and the end devices, STP operation should be avoided due to the active/passive redundancy. However, STP does offer advantages in some scenarios:

- Where construction-related restrictions limit the number of possible connections, STP is the ideal alternative. This minimizes the risk of loops forming, especially in client-access mode.
- With its modest hardware requirements the protocol can be supported even by entry-level switches, which makes STP a very cost-effective solution.

The supporting protocols LACP, VRRP, DHCP relay, and L3 routing

In addition to the three protocols already mentioned, which significantly determine the overall concept of the switch network, further protocols are important for the following scenario description.

Link Aggregation Group (LAG) & Link Aggregation Control Protocol (LACP)

The technology for implementing link aggregation and load balancing is called LAG (Link Aggregation Group). An LAG dynamically bundles a number of physical connections between network devices into a single logical connection.

LACP is the acronym for "Link Aggregation Control Protocol". As part of the global standard IEEE 802.1AX (Link Aggregation), LACP is a protocol for the automatic configuration and maintenance of link aggregation groups. LACP uses LACPDU (LACP data packets, request-response principle) as an automated negotiation mechanism between two or

when using VPC or stacking, several network devices, so that a logically grouped link can be automatically formed and started according to its configuration. LACP is also responsible for maintaining the link status and constantly exchanging information about the data packets. It therefore reacts dynamically to changes in the network without the need for reconfiguration.

In addition to the redundancy provided by two independent physical connections, the LACP protocol associated with the LAG not only shares the load between the connections it operates, but uses LACPDUs to automatically detect whether a connection has failed or is being reconnected without any detectable data loss, so maximizing throughput. This last aspect offers a major advantage over STP, which only uses one of the two physical connections, with the other one only ever being used for connection establishment.

Virtual Router Redundancy Protocol (VRRP)

VRRP is a standardized layer-3 network protocol that uses redundancy and load balancing to provide automatic allocation and dynamic failover to keep routers available, or in this case switches that support routing. This ensures network availability, especially for security-critical services, through the seamless transition to a backup device. In very large networks (campuses with more than 10,000 ports), the routing concept required on layer 3 can also be simplified, as the two devices in the VRRP can be virtualized as a single default gateway.

DHCP relay

Since two-tier or three-tier networks usually have a separate DHCP server on high-performance hardware, it is important for switches on the aggregation/distribution and access layers to be configured with a DHCP relay agent. This forwards DHCP requests to a centralized DHCP server and prevents IP-address conflicts.

Layer-3 routing

Routing functions are essential for implementing security and the options of access control, dynamic growth of the network and good stability (forwarding vs. flooding) via a logical and above all efficient separation of subnets. To ensure that each switch knows which router to use, a routing table is created that serves as an "address database" that is valid at all times. Dynamic routing ensures that all "routers", i.e. layer-3 capable switches (L3), can communicate with one another and build this routing table

independently. This means that the route of data traffic within the network is constantly being set dynamically, which ensures the best network performance. Common routing methods are OSPFv2/v3 and BGP4, although the former is generally used only in internal networks.

Example scenarios for redundant switch networks

Now that we are familiar with the protocols and their core function, we now move on to their application in example scenarios with models from the [LANCOM switch portfolio](#).

The examples shown deal with three-tier switch networks. If a two-tier network with aggregation/distribution and access layers is sufficient for you, the core layer can be omitted. The solutions described remain valid and can be seen as recommendations for practical application.

Scenario 1: 100%-uptime switch network with VPC-capable access switches

This scenario is suitable for large enterprise and campus networks with high redundancy requirements. The maximum number of access ports with 100% redundancy is approx. **60,000**.



In the case of a core switch with 32 ports, one port is usually used for the uplink, e.g. to a data center/WAN, and another 2 to 8 are reserved for VPC offering redundancy and performance. So with 6 VPC connections, 25 ports remain. On the aggregation/distribution layer, redundant switches with 48 ports each are connected. In turn these can connect to switches on the access layer, each with a maximum of 48 ports. This results in

$$25 \times 48 \times 48 = \mathbf{57,600} \text{ ports}$$

To implement this scenario, all switches from the core to the access layer must be VPC-capable. Although this limits the number of possible switches, it ensures a high bandwidth in the active/active principle with 100% uptime and in-service software upgrades (ISSU) for the highest network requirements.

This scenario is ideal for the most powerful LANCOM switches, such as the core switch [LANCOM CS-8132F](#), the aggregation/distribution switch [LANCOM YS-7154CF](#) as well as the [XS-4500 series](#) access switches. For the first time, the XS-4500 series enables the connection of Wi-Fi 7-capable access points such as the [LANCOM LX-7500](#).




Figure 2:
Scenario using VPC-capable
access switches

The switches at every network layer are connected via 100G VPC peer links. The lower layers are then redundantly connected via LAG with 100G or 25G, depending on the uplink ports of the access switches. It can also be seen that core-layer switches in the VPC group are configured with VRRP. This helps to simplify the subsequent routing configuration on the lower layers, since VPC-enabled switches keep their respective IP addresses and it is only VRRP that then simplifies these down to a single shared one. Consequently the switches at the core and aggregation/distribution layers appear from the access layer to be a single L3 routing gateway. Not shown are the auxiliary protocols DHCP relay and dynamic routing such as OSPF. These should be configured and used according to their intended function in order to make network segmentation with VLANs as simple as possible.

For an end device (e.g. access points) to be fully redundantly connected and allow for a meaningful VPC configuration, it requires at least two Ethernet ports. This is because a VPC connection is always based on redundancy and uses a bundled connection (LAG). If an access point has two data ports, the power supply remains uninterrupted thanks to the non-stop PoE function of LANCOM access switches, even when one of the VPC switches is restarted or updated.

Devices with only one Ethernet port can also be connected to a VPC layer using a static LAG, but this comes with several disadvantages: In this case, the second port is only created virtually, causing a physical port on the second VPC switch to be blocked and remain unused. Furthermore, this setup allows for a maximum of 48 access points to be connected ($2 \times 48 / 2$). Therefore, this workaround is not recommended by LANCOM and is not configurable for certain devices, such as printers or PCs.

Scenario 2: Reliable switch network with a combination of VPC and stacking

This scenario focuses on the costs per port. If it is possible for the access layer to work with maintenance windows, this scenario with stacking at the access layer is the recommended method. In contrast to the first scenario, the aggregation/distribution layer here can operate for example the LANCOM XS-6128QF, and the access layer can operate the more cost-effective GS-4500 instead of the XS-4500 series. Since it is now possible to plan with up to eight switches in the stack on the access layer, the number of ports increases to a maximum of **460,800** ports (25*48*48*8). This significantly increases the number of ports while maintaining an acceptable level of redundancy and close to 100% network uptime (assuming there is a maintenance window).

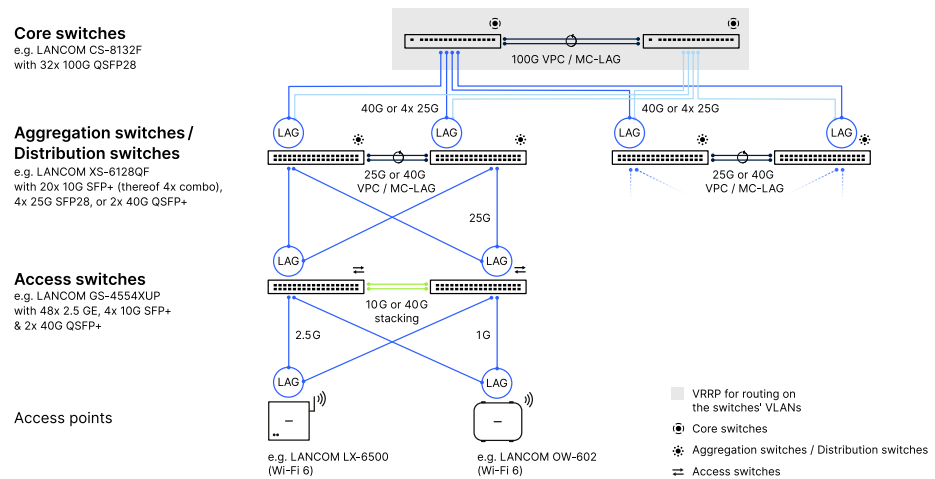


Figure 3:
Scenario with a combination of
VPC and stacking

Due to the high number of ports, the use of the L3 routing protocols VRRP and ARF (Advanced Routing and Forwarding) should be considered at core level. VPC remains in place at both the core and aggregation / distribution level and, as in the first scenario, fulfills the requirements of the ISSU approach.

At access level, the redundancy solution stacking is used instead of VPC. This enables the use of a larger number of access switches from the LANCOM portfolio. DHCP relay and the use of LAGs between the levels remain the same as in the first scenario.

However, stacking comes with certain limitations: During a firmware update of the switch stack, a downtime of approximately five minutes is to be expected, which makes it necessary to plan a maintenance time window. One key advantage of this scenario, however, is that a significantly larger number of end devices can be connected without any problems due to the absence of VPC and LAG redundancy restrictions.

Scenario 3: Cost-optimized switch network with a combination of VPC and STP

In this scenario, the configuration of the core and aggregation / distribution layer with VPC and LAG remains unchanged as before. Only the LANCOM access switch LANCOM GS-3652XUP ensures diverging uplink speeds.

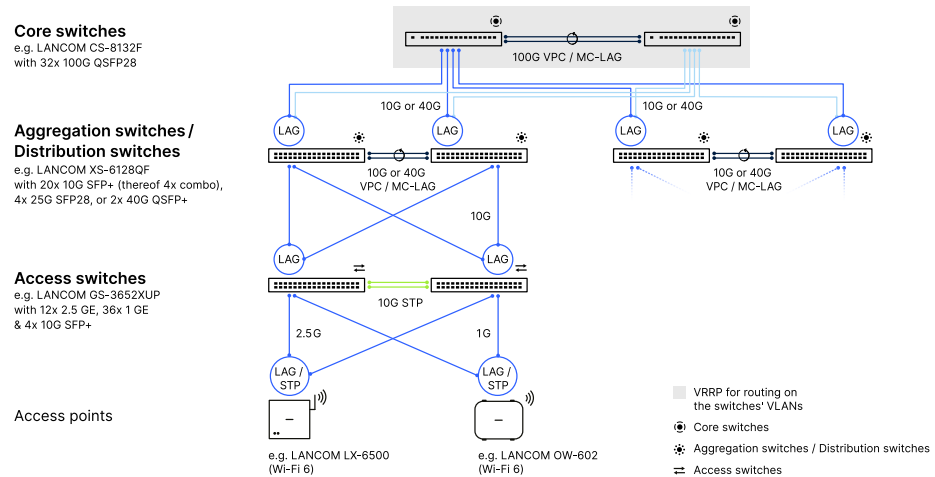


Figure 4:
Scenario with a combination of VPC and STP

On the access layer, STP is configured instead of VPC or stacking. This has the advantage that the protocol requires only modest hardware performance, which further increases the selection of viable access switches (e.g. the [LANCOM GS-3600 series](#)). However, STP has only a limited range of uses due to the active/passive principle and the laborious configuration.

In the following, we present two typical examples to illustrate the use of STP.

Scenario 3.1: STP at decentral sites

The two aggregation/distribution switch stacks should be understood to be two independent units at different locations. Using LACP and the STP configured on it, both stacks are now connected to the backbone which also contains the gateway to the WAN. If the connection from the right-hand stack to the WAN gateway fails—for example, due to unforeseen events—the stack can still route to the WAN via the left-hand stack without the site being cut off completely. As long as there is no error, the middle connection between the stacks stays inactive. On the access layer, the recommendation for this scenario is still to use LACP instead of STP.




Figure 5:
STP at decentral sites

Scenario 3.2: STP with numerous cascaded access switches

This scenario is ideal when the budget is limited but a large number of access ports still needs to be implemented. Cost cutting often targets the stack of aggregation switches because there is no avoiding the large number of access switches. To retain a certain amount of redundancy, a ring is configured on the access layer, which requires the activation of STP. It is also possible to set up double connections via LACP here. However, this can also be omitted here due to the cost aspect.

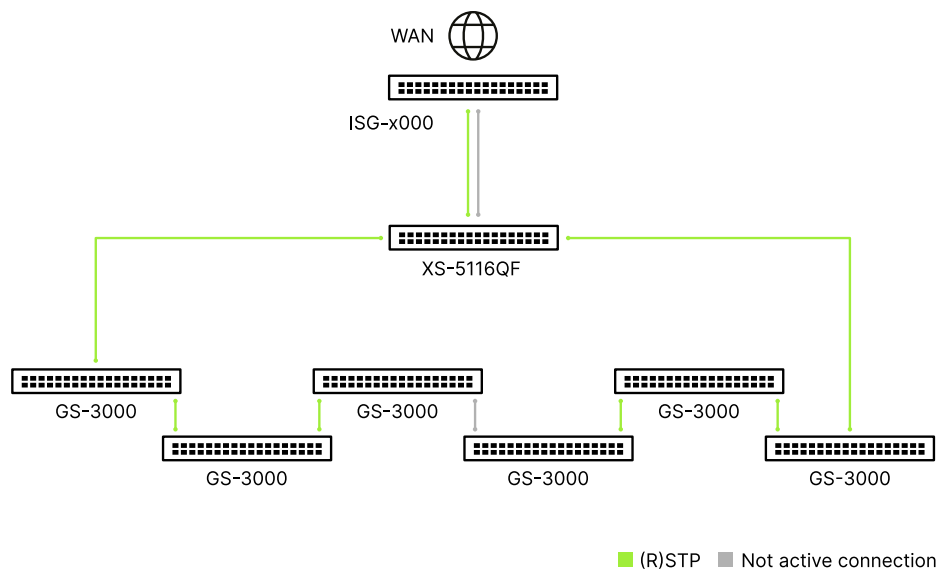


Figure 6:
STP with numerous cascaded
access switches

Conclusion

By expanding their portfolio to include the core layer, LANCOM has become a one-stop shop for anybody planning or managing campus networks.

Even if these scenarios cannot reflect every possible network design, these examples give a good overview of what can be achieved with LANCOM core-, aggregation/distribution-, and access switches. With the redundancy concepts VPC, stacking, and STP presented here, the best solution for any network requirement can be found depending on the application and budget.

Are you planning to set up or expand your network with LANCOM switches?

Experienced LANCOM technicians and the specialists from our system partners will help you with the planning, installation and operation of a needs-based, high-performance and future-proof LANCOM network designs.

Do you have any questions about our switches, or are you looking for a LANCOM sales partner?

Please give us a call:

Sales in Germany

+49 (0)2405 49936 333 (D)

+49 (0)2405 49936 122 (AT, CH)

