

7 reasons to stack your core chassis

Why use stacking?

Every network needs a resilient and versatile core. A chassis can seem like a great solution, since it allows double the controller cards and power supplies for resiliency, and it can use any combination of line cards for versatility of speeds and feeds. However, with a chassis-based solution, it all has to be located in one place. If the network is using a redundancy protocol like Virtual Router Redundancy Protocol (VRRP), one chassis is lying in standby not being used, which is expensive. Plus, the failover performance of VRRP means that downstream hosts will see a noticeable outage in the event of a failure.

There is a lot of information about stacking and its benefits, such as easier management and fast recovery from failures. But a stack of individual fixed-format devices does not yield the flexibility that is needed to change or grow speeds and feeds easily. This type of stacking also limits the number of available ports. A stack of chassis is far more useful. It offers all the benefits of a chassis, but can be located in different geographical locations to offer protection from local disasters. It also offers very fast failover performance so that downstream hosts do not notice any failures. With chassis stacking, full use can be made of both chassis, because it is an active/active architecture. Investment on equipment is better utilized, and available bandwidth is doubled.

But is stacking really necessary? Multi-Chassis Link Aggregation (MLAG) works well in the data center to provide a dual-homed aggregated link for servers. It is active/active and has fast failover without complex protocols. Can MLAG be used for core chassis too? The answer is yes, but if the network also has to forward Layer 3 traffic, the complexity of the MLAG solution increases dramatically. This is because MLAG is best-suited to Layer 2 applications, like data centers. For Layer 3 networks, a single MAC address and a single IP address are required to minimize downstream disruption after a failover, and a method of failing over the Layer 3 routes is required. The standard method for creating a redundant unicast Layer 3 core is to use Virtual Router Redundancy Protocol (VRRP) to provide a gateway failover for hosts in the local LAN. However, VRRP does not provide a rapid failover, and also puts one chassis into standby mode. This solution satisfies none of the uptime, energy efficiency or ROI expectations that organizations now have of their data networks.

For multicast, there simply is no standard protocol for managing router redundancy. Stacking takes care of the dual-homed aggregated links and provides a simple solution for a virtual IP address, so that Layer 3 traffic can be routed without the need for complex failover protocols. Additionally, the stack is still fully active/active even for Layer 3 traffic, since the states of Layer 3 protocols are automatically synchronized across both stack members, which leads to minimal traffic disruption in the event of a failover.

VCStack Plus

Allied Telesis VCStack Plus allows two SwitchBlade x8100 chassis to be stacked with either one or two SBx81CFC960 controller cards in each, for easy active/active resiliency with rapid fault recovery. Initial configuration is effortless, requiring no spanning tree or failover protocols, and ongoing administration is reduced because both chassis can be managed as a single virtual device.

Allied Telesis designed VCStack Plus to meet the following seven requirements of a resilient network core solution:

- 1. Fast failover:** no downstream disruption on failover; predictable failover times
- 2. Simple configuration:** reduced maintenance and less chance of configuration errors
- 3. Maximizing return on investment:** active/active architecture for reduced equipment costs
- 4. Cross-chassis LAGs:** ideal for resiliency, bandwidth and easy management
- 5. Scalable:** ability to add more ports without having to reconfigure or replace equipment
- 6. Powerful simplicity:** the benefits of multiple stacked units with ease of management
- 7. Long distance:** perfect for distributed network environments and protection against local disasters

1. Fast failover

A key requirement of any resilient system is its ability to recover quickly from failure. For a network core this is especially important since by definition, the effects of disruption in the core are magnified in the outer network layers.

Traditional failover mechanisms rely on standby devices monitoring the health of active devices using CPU-intensive protocols. When the standby device detects that the active device has failed, it takes over as the active device. This can create a significant disruption to the attached hosts, because the standby device can take up to 30 seconds to detect that the active device has failed, depending on the failover protocol being used. Faster protocols are not the answer, as they consume more CPU cycles and make the entire system less stable. During failover, tables may also be flushed and addresses re-learned—causing more disruption downstream.

Stacking uses a different failover mechanism because it keeps both devices in sync with each other. If one fails, another takes over seamlessly with no disruption to attached hosts. This is because a stack is tightly coupled, so failure detection is almost instant (sub-second); and because tables do not need to be flushed. Failover times are also predictable in a stack because the time to detect the failure is consistent, unlike protocols which can have a wide variance.

With stacking, failover is fast and the user experience is predictable. Network administrators can perform maintenance tasks during the day, safe in the knowledge that users will not notice any disruption to their traffic.

2. Simple configuration

Studies and practical experience shows that one of the most common causes of failure in a network is configuration error. The easier a configuration is to understand, the less likely it is that mistakes will be made, and the easier it will be to troubleshoot and diagnose issues.

Failover protocols, such as Hot Standby Router Protocol (HSRP) and VRRP, are typically used in combination with Spanning Tree Protocol (STP) to provide automatic switchover to the standby system when the active system fails. But these protocols require significant configuration, and diagnosing issues can be time-consuming because of the complex interactions and timing dependencies between processes.

Stacking offers a much simpler way to configure a resilient pair of chassis. There is no need for failover protocols or STP, and the configuration commands are simple. Diagnosing issues is much easier too—stacking has in-built diagnostics and generates Simple Network Management (SNMP) traps and log messages whenever a significant event occurs.

Because stacking virtualizes the control plane across multiple devices, these devices operate as a single unit under unified software control. The configuration that is applied to the stack does not have to specify the way that information is communicated within the stack. It simply treats the stack as a single virtual device, and needs only to specify how the stack interacts with its external environment. The software comprising the virtualized control plane manages the route distribution, bandwidth sharing, forwarding plane synchronization and more within the stack. This centralized control of a distributed set of devices is an effective way to simplify network core management.


Another benefit of stacking is that the effort required to manage both chassis is halved because stacking presents them as a single entity. There is only one management interface, so stack management is as simple as managing a single device.

Both stacked devices run the same configuration file, so changes only need to be made once because they are synchronized on both devices automatically. This reduces the chance of errors and saves time. Firmware versions are also synchronized across stacked devices, so firmware upgrades are easier too.

3. Maximizing return on investment

A multi-chassis network core represents a major investment for an organization. It is not only a major investment of capital expenditure (CAPEX), it also requires significant ongoing operational expenditure (OPEX) in the form of IT engineering time, electricity usage, upgrade expenses, support contracts, and more. Therefore, an organization needs to maximize the benefit it gains from this investment. If half of the network core sits idle as a standby redundant unit, this is not an effective use of this investment.

Unfortunately, this is the typical situation when traditional failover protocols are used with STP. One device is designated as the active system and performs all the control and data plane processing. The other device waits in a standby mode, monitoring the active system and ready to take over if a failure is detected. The resulting redundancy is therefore very expensive, and the network cannot take advantage of the potential bandwidth sitting idle in standby.



The above diagram shows how STP blocks links to prevent loops, but causes reduced bandwidth.


Stacked systems can maximize the available bandwidth because they operate in an active/active mode, and not the active/standby mode of other solutions. One member of the stack is nominated as the 'Master' and manages the control plane traffic for the entire system. Crucially however, all stacked devices can forward data plane traffic, because they are all active forwarding engines. STP is not used, so none of the links to the stack are blocked. The maximum bandwidth is available to be utilized.

If a stack member fails, the available bandwidth is reduced to that of a single device until the fault is repaired. However with stacking, reduced bandwidth is the exception—whereas with failover protocols, it is the norm.

4. Cross-chassis Link Aggregation Groups (LAGs)

A Link Aggregation Group (LAG) allows bandwidth to be increased on a link, by adding multiple physical connections together into a logical link. LAGs are easy to configure and require minimal management. LAGs also provide link resiliency, because if a physical connection in the LAG fails, then traffic is automatically and instantly redirected to the remaining connections. LAGs are commonly used as connections to servers with teaming Network Interface Controllers (NICs) where the benefits of increased bandwidth, rapid fault recovery and easy configuration are very valuable.

The benefits of LAGs mean that they seem like an ideal way to connect devices to the network core. Unfortunately, a LAG can only be configured between two devices, not three. So it is not possible to use LAGs with traditional failover protocols to spread the connections across both the core chassis and the downstream device.



The above diagram shows how a stacked system allows all links to remain active. LAGs can be configured across both chassis because they belong to the stack and are treated as a single virtual device. Bandwidth is maximized and recovery from link failures is almost instantaneous. For time-critical applications, using LAGs helps to reduce latency—traffic can always follow the most direct path because no links are blocked.

VCStack Plus with LAGs provides a system that makes full use of the network, and has no single point of failure.

5. Scalability

Typical chassis-based systems offer a number of line card slots that may be used with a wide variety of line cards offering different speeds, feeds and functions. Stacking doubles the available number of slots without increasing the administration overhead, thereby offering much greater scalability and flexibility than a single chassis.

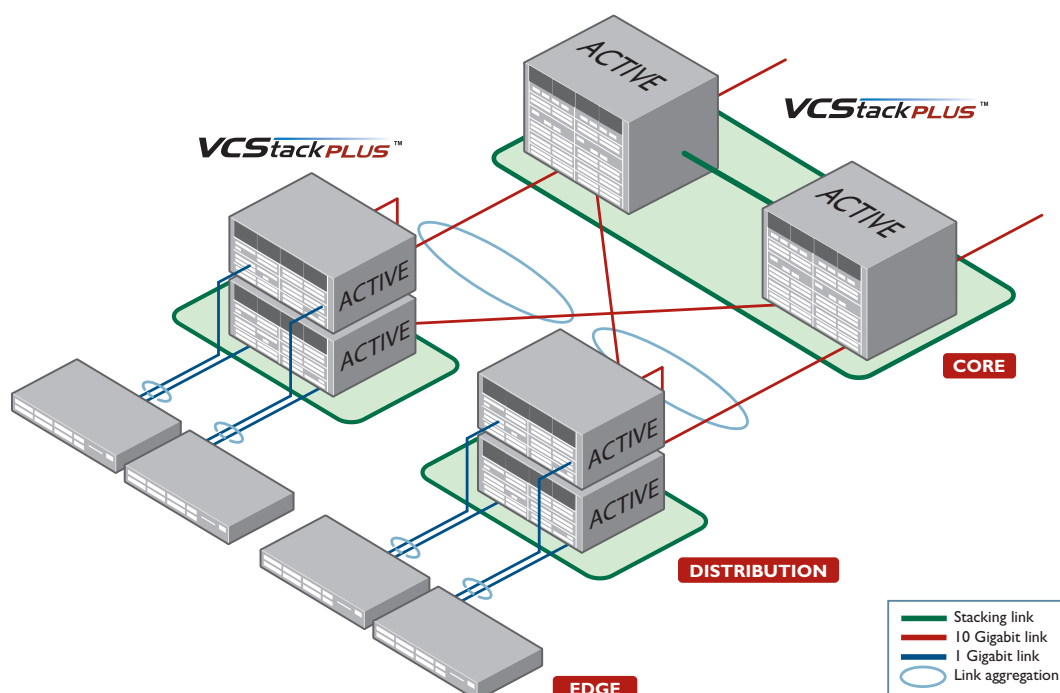
As the stack operates as a single entity within a virtualized control plane, adding and changing line cards creates low overhead cost, from the configuration management point of view. Once a card is hot-swapped into one of the chassis, it is immediately visible to the whole stack. Synchronization of forwarding information with the new hardware occurs automatically. The presence of the new hardware is immediately taken into account by the route and bandwidth distribution processes within the stack. No user intervention is required to ensure that routes, multicast groups, IPV6 neighbour entries, etc. that are learnt by this newly installed hardware are automatically shared with the whole distributed core, and vice versa.

6. Powerful simplicity

In larger networks, multiple stacked chassis can also be employed at the distribution layer for increased resiliency and greater port density per location.

If multiple separate chassis were used in the network core and at distribution layer locations, the network would become extremely complex to manage and configure. Redundancy and spanning tree protocols would be required for unit standby and loop protection. With a number of chassis sitting in standby, network bandwidth would be wasted, and investment return in network infrastructure not fully realized. Furthermore, failover performance would be suboptimal.

Using VCStack Plus in both the core and distribution layers of the network means that each location is seen as a single virtual unit, and the full network power is utilized. With fully redundant pairs of appropriately-sized chassis at each location, and LAGs connecting them together, this solution scales to meet the requirements of very large networks, while remaining easy to manage - powerful simplicity.




7. Long distance

The increased distance provided by fiber stacking connectivity means that VCStack Plus members do not have to be co-located, but can be kilometres apart. This enables a solution that is perfect for distributed network environments.

Two different buildings on a campus can both be part of the network core, with the full redundancy of a chassis provided in each location, incorporating dual hot-swappable power supplies and control cards to maximize uptime. Both distributed chassis retain all of the benefits of VCStack Plus, and LAGs from downstream devices can be spread across the units for complete resiliency. The VCStack Plus network core continues to function as a single entity for simplified cohesive management.

Long distance VCStack Plus can also be employed to provide a data mirroring solution that retains the simple management and resiliency of a single network, while providing protection from local disasters.



Summary

Chassis-based systems have a huge advantage over fixed-format devices as they can be easily expanded with additional line cards and modules. Speeds and feeds can be changed and increased to cope with business demands, without major reconfiguration of the network or a need for more rack space. Even a stack of fixed-format devices is less flexible than a chassis because chassis line cards tend to have more interface options, are easier to install, and are usually cheaper than similar fixed-format devices.

However chassis-based systems are more complex to configure and manage than a fixed-format device, and require specialist administration skills which are expensive. So many companies select a stack of fixed-format devices instead of a chassis for their network core because the on-going operating expenses are lower.

Chassis stacking with Allied Telesis VCStack Plus offers the best of both worlds—superior flexibility of speeds and feeds, and simplicity of management leading to lower operating costs. A chassis stack is a more attractive solution for cost-conscious businesses than a stack of fixed format devices.

Compared to traditional core network failover mechanisms, VCStack Plus chassis stacking provides:

- increased reliability of the network core because there are two devices operating as one
- more cost-effective use of equipment because both chassis are active
- increased bandwidth utilization because no links are blocked
- faster recovery from link or device failures
- simplified configuration and management leading to lower operating costs and fewer mistakes

Allied Telesis VCStack Plus provides the complete solution for resilient chassis-based network cores.

About Allied Telesis, Inc.

For nearly 30 years, Allied Telesis has been delivering reliable, intelligent connectivity for everything from enterprise organizations to complex, critical infrastructure projects around the globe.

In a world moving toward Smart Cities and the Internet of Things, networks must evolve rapidly to meet new challenges.

Allied Telesis smart technologies, such as Allied Telesis Autonomous Management Framework™ (AMF) and Enterprise SDN, ensure that network evolution can keep pace, and deliver efficient and secure solutions for people, organizations, and “things”—both now and into the future.

Allied Telesis is recognized for innovating the way in which services and applications are delivered and managed, resulting in increased value and lower operating costs.

Visit us online at alliedtelesis.com