White Paper

Allied Telesis

Distributed Collapsed Backbone

How a network can have its cake and eat it too.

Introduction

Network backbones come in two broad varieties: “distributed” and “collapsed”.

A **distributed backbone** has a core consisting of multiple switches or routers chained together, typically in a ring.

A **collapsed backbone** has a central device at the hub of a star network. In medium to large networks, this central device is a chassis switch.

Each of these two designs has its advantages and disadvantages, and some environments are better suited to one design in particular. However, there are some network environments which could utilize the benefits of both options combined– particularly if they could avoid the disadvantages of both options.

This white paper discusses how to create a network design that combines both distributed and collapsed backbones. This combined design is called a **distributed collapsed backbone**.

To understand how such a design works, it is necessary to first consider some of the characteristics of these two designs, to understand their benefits and downsides, and to understand the environments to which each is most suited.

Distributed backbones

The standard model of a **distributed backbone** network consists of a core ring of routing/switching nodes, each of which is the hub of its own local star network or its own local ring, which then extends out to a local access layer.

In some cases, there may be servers connected locally to most of the backbone nodes. In other cases, the servers are all located at a single node at the 'head end' of the distributed backbone.

In either case, traffic from the users local to each node enters the backbone via that node, and at least some of the traffic traverses through other backbone nodes to reach its destination. Thus, the links within the backbone carry traffic to and from destinations throughout the network, not simply to and from destinations that are local to the nodes at each end of the link.

A distributed backbone design is especially suited to physically spread out environments, for example a campus network that connects distinct departments into a unified network.

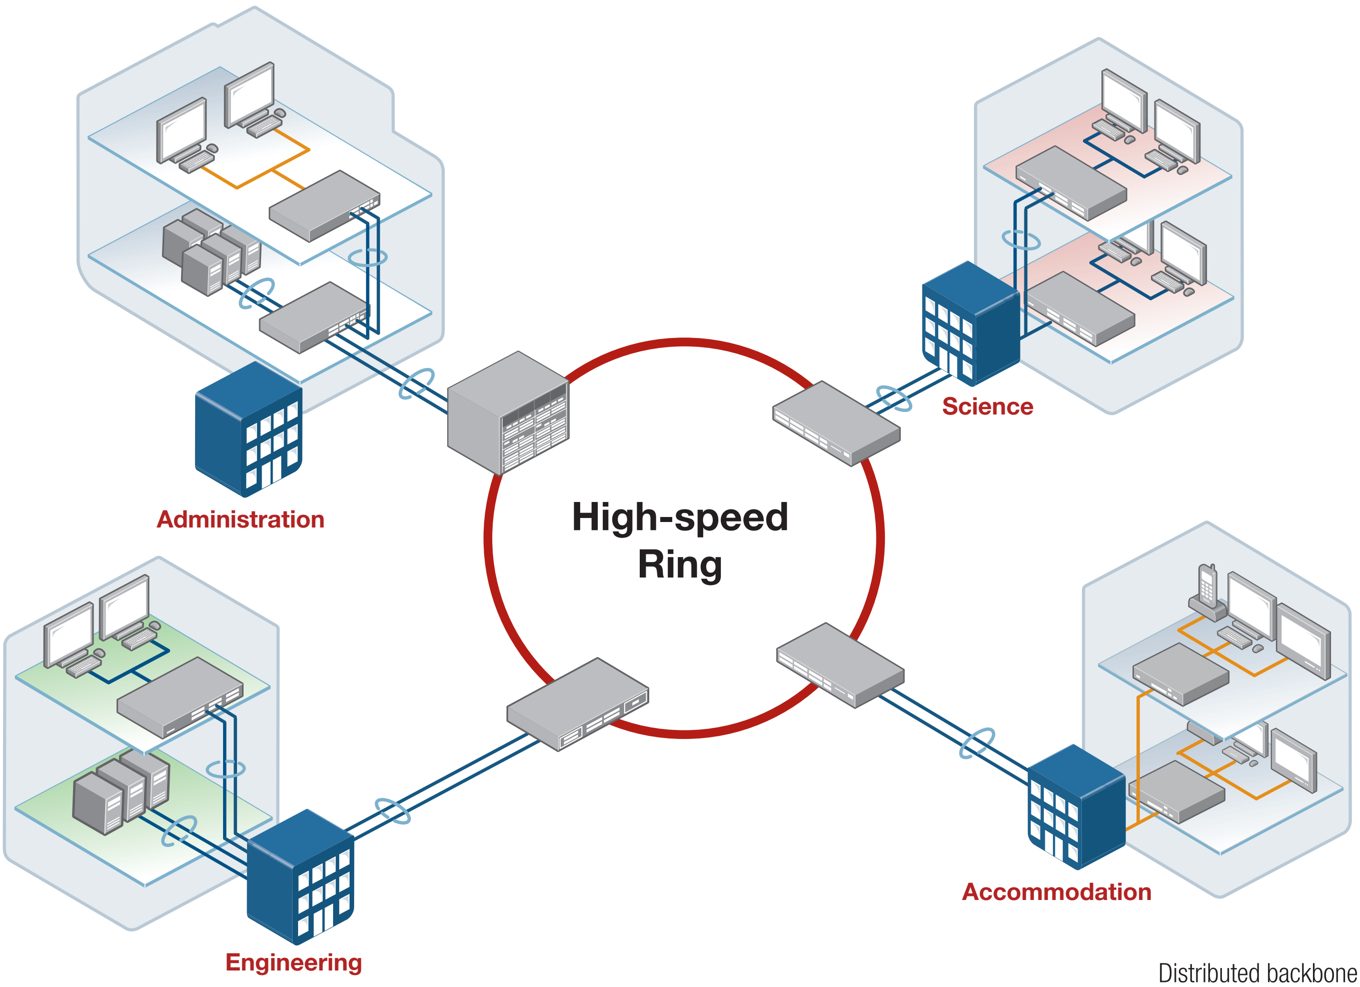

Distributed backbone

The distributed backbone network design, like the collapsed backbone network design, has distinct advantages and disadvantages, as follows.

Advantages

Reduced failure impact

The primary advantage of a distributed backbone is that it can reduce the impact of the failure of any given node. This is particularly true if services and resources are distributed around the nodes. If the backbone is in the form of a ring, then the failure of a single link in the ring does not cut users off from services, as their traffic will still have an alternative path to reach any location within the network.

In the ring topology, even the failure of a backbone node will only disrupt the users and services connected to that node–all other users will still be able to access all other services.

Network segmentation

The distributed backbone also has the effect of segmenting the network, which can be advantageous in localizing the impact of network problems, and aiding in the troubleshooting of those problems. This is particularly so if the backbone nodes are performing Layer 3 routing between their local users and the backbone. In this case, the effects of any loop-induced storms or link flapping are isolated to that local region, and not propagated further into the network.

Disadvantages

Backbone links are shared by many users

The prime disadvantage of a distributed backbone is that the backbone links are shared by users and services located right around the network. If there is a growth in the bandwidth needs of the users who are local to one node of the backbone, that growth can potentially affect multiple backbone links. This is particularly so if most of the services these users wish to access are multiple backbone hops away from their own location.

Difficult to balance bandwidth allocation

Allocating fair portions of backbone bandwidth to different portions of the network can be an ongoing balancing act, requiring regular subtle adjustment as requirements evolve.

Single point of failure is not entirely eliminated

If the bulk of services and resources are local to a single head-end backbone node, then this node represents a single point of failure. Hence, the distributed nature of the backbone is of limited advantage in the case that most of the network operation depends on access to a single location.

Collapsed backbones

The term **collapsed backbone** refers to a network design in which the nodes that create the network backbone are "collapsed" or compressed together into a single core device. It creates an image of the core ring collapsing, until it has all consolidated at a single point in the center.

In recent years, the collapsed backbone has become almost synonymous with the chassis switch. A powerful multi-blade switch is a compelling choice for the central device to sit at the core of a medium or large network.

The collapsed backbone design is well suited to enterprise networks, such as hospitals, government, schools and general business environments. It is suited to both networks contained in a single building, and to those spread out over a moderately sized area. The design is not as well suited to a more physically spread out environment, for example a metro service network, where individually cabling numerous remote sites back to a core can be cost-prohibitive.

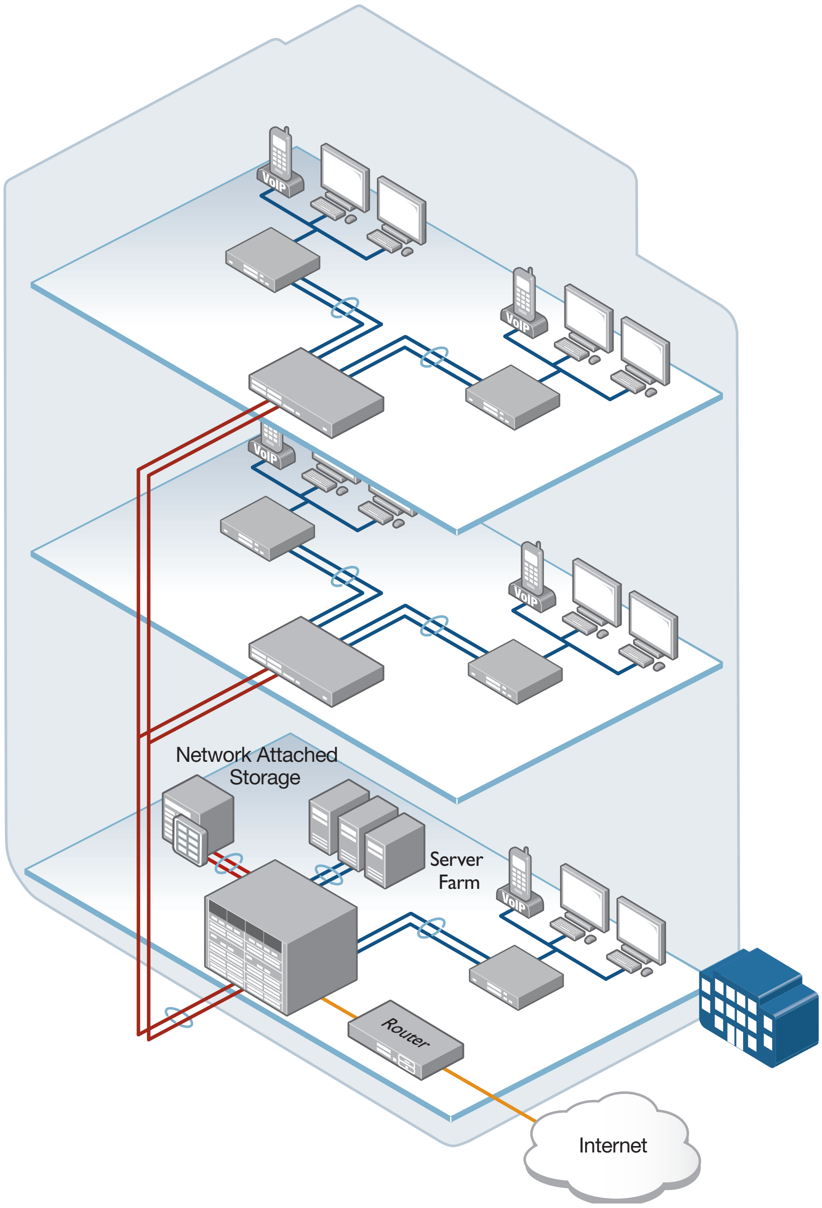

Collapsed backbone

Advantages

Simple bandwidth allocation

Collapsed backbone networks use a star topology, which greatly facilitates the provisioning of appropriate bandwidth to the connected nodes. Each connected distribution node has its own dedicated connection into the core. There is no multiplexing of different nodes' data onto shared backbone links, as occurs in a distributed backbone.

As a result:

- Bursts of traffic from one distribution node in the network do not directly impinge on data forwarding from other nodes.
- It is straightforward to monitor the bandwidth requirements of any given distribution node, as it is simply a matter of measuring the usage on the dedicated link out to the node.
- Increasing the bandwidth provisioned to a given node is uncomplicated–it is just a matter of upgrading the link to that node. In contrast, upgrading capacity to a node connected to a distributed backbone may involve upgrading multiple backbone links.

Reduced Capex and Opex

A collapsed backbone design means that the investment in complexity and high-performance is focused on just one device. This has both Capex and Opex advantages:

- **Capex - investment concentrated on one core node**
  The investment in advanced capability does not have to be multiplied across multiple backbone units. Instead, it can be dedicated to just the one core node. This has two potential Capex advantages. It can mean a cost reduction, due to the need to buy high-capability equipment for just one location. It can also mean a significant step up in technology–the money that would have been invested in multiple moderately capable backbone units, can instead be utilized to create one very high capability node.
- **Opex - simplified configuration and troubleshooting**
  The advantage lies in simpler configuration, and less diversity in troubleshooting activity. Distributed backbone networks involve protocol interactions between backbone nodes, so it is necessary to configure and maintain resiliency and routing protocols among a set of backbone units. In a collapsed backbone, there is only one backbone unit, so there is no requirement for protocol interactions within the backbone. This greatly simplifies configuration requirements.

  Further, in a collapsed backbone network, the configuration of backbone security, Quality of Service (QoS), Virtual Local Area Networks (VLANs), multicast routing, and other skill-intensive features do not have to be distributed across multiple units. The configuration of these features simply needs to be accurately completed on one unit, and then kept up-to-date on that one unit.

  Further, troubleshooting activities can be reduced by proactive monitoring of the one core node–rather than diluting monitoring efforts across multiple backbone units.

Extra advantages when a chassis switch is the core device

The benefits of the collapsed backbone design are magnified when the core device is a chassis switch. A chassis-based core can supply periodic technology refreshes with little network disruption, and affordable incremental cost.

The advantages that a chassis switch provides are:

- **Internal system redundancy and resiliency.**
  A chassis provides redundant PSUs, redundant CPUs, redundant cooling capacity and more. Additionally, if the connections from the core chassis to distribution nodes consist of aggregated links, terminated on different blades of the chassis, then the network is also resilient to the failure of any one blade.
- **The port count of the chassis can be increased incrementally.**
  As the requirement for links into the core grows, new blades can be purchased and hotswapped into the chassis to effect a simple, non-disruptive increase in port density.
- **Chassis switches invariably provide a hardware and software roadmap.**
  Over time, older blades can be swapped out, and replaced with newer-generation blades that provide higher-bandwidth ports, larger forwarding tables, better QoS and improved security. Software upgrades unlock these extended hardware features, as well as bringing in new diagnostics, advanced management capabilities, and conformance to protocol revision updates.

Disadvantages

One single point of failure

The primary disadvantage of the collapsed backbone is that it represents a single point of failure. Although module redundancy–for example CPU, PSU, and blade redundancy–within a core chassis mitigates the risk of a complete core outage, the fact remains that a backbone consisting of a single core node, at a single location, is more at risk of a complete backbone outage than is the case for a distributed backbone.

Increased installation costs

More cabling is required in a collapsed backbone network, as each distribution node requires its own connection to the core. This can increase the initial installation cost of a collapsed backbone network, compared to a distributed backbone network.

As well as this, the cabling required to add new distribution nodes to the network is typically more in a collapsed backbone than in a distributed backbone network. In the distributed backbone, a new distribution node may have just a short distance to connect to its local backbone node. In a collapsed backbone, every new node must connect right back to the core node.

Getting the best of both worlds: the "distributed collapsed backbone"

Each of the distributed and collapsed backbone network designs has its own set of advantages and disadvantages. The ideal would be a new network design–one that combines the two designs to feature their main advantages, and avoid the main disadvantages.

Such a design is now possible. The answer is to design and build a network backbone that consists of a cluster of individual units operating as a single logical unit.

This combination backbone is both distributed and collapsed:

- **Distributed**
  It consists of multiple units, which can be placed at distinct physical locations.
- **Collapsed**
  The backbone operates as a single unit, sharing a single control plane that automatically synchronizes forwarding information across the cluster, presents a unified management interface, and coordinates the collection of diagnostic and statistical information.

It is literally a **distributed collapsed backbone**.

Advantages

A backbone formed in this way features the main advantages of both the collapsed and distributed backbones:

- No single point of failure–the cluster will continue to operate even if one or more of the clustered units become unavailable.
- Simplified management and troubleshooting–the cluster is managed, and monitored, as a single unit.
- Simplified configuration–no routing or resiliency protocols need to be configured between the clustered units. The control plane that is distributed across the cluster takes care of sharing forwarding tables between the units, and of rapidly recovering from the loss of any intra-cluster links.
- Adding new nodes into the core cluster is simple. A minimally configured unit, cabled appropriately into the cluster, will be automatically absorbed into the cluster and provisioned with the correct configuration. It will then begin immediately to operate as a fully functional member of the cluster.
- Software upgrades can be applied to the whole cluster as a single action, in a near-hitless fashion.
- If the cluster is simply a clustered pair of chassis switches, then this design can utilize all the benefits inherent in chassis switches.
- Cabling from distribution nodes to the backbone does not need to extend all the way to a single core location, but can terminate at a local cluster node.

Disadvantage

The distributed collapsed backbone design retains only one of the disadvantages of the distributed backbone or collapsed backbone designs. There is still sharing of different nodes' traffic across the backbone links. The traffic from users local to different cluster nodes will be multiplexed onto the intra-cluster links, and these will then compete with each other for available bandwidth.

The Allied Telesis distributed collapsed backbone offerings: VCStack™ and VCStack Plus™

The Allied Telesis switch clustering technology is called Virtual Chassis Stacking (VCStack). The technology is implemented in two closely related, but slightly different, forms:

VCStack

The stacking of Allied Telesis pizza-box and SwitchBlade® x908 Generation 2 switches provides excellent options that are optimized for different size networks.

The option most suitable for implementing a distributed collapsed backbone is the Long-Distance Virtual Chassis Stacking (VCStack-LD)—enabled by using fiber stacking links. This solution provides a backbone of multiple high-performance switches which can be distributed over distances up to tens of miles.

VCStack Plus

A pair of SwitchBlade® x8100 Series chassis switches can be combined into a single logical unit. The controller cards in the two chassis operate in unison to provide complete synchronization of all Layer 2 and Layer 3 tables for IPv4 and IPv6 unicast and multicast forwarding. The stack can coordinate up to 4 controller cards—2 per chassis—and is resilient to the loss of any controller card, provided any one controller is still in operation.

Like VCStack-LD, the two chassis in the stack can be separated by up to tens of miles.

The VCStack Plus solution represents an ideal means of distributing a chassis-based collapsed backbone across two physically separate sites to provide a data-mirroring and automated disaster recovery solution.

If cabling capacity allows, an excellent level of resilience, and traffic sharing, can be achieved by connecting distribution nodes to the VCStack Plus core via aggregated links that combine ports on both chassis as shown in the diagram. This combination creates a backbone with no single point of failure whatsoever. PSUs, controllers, blades, inter-chassis links, distribution-to-core links—even whole buildings—can be lost and users remain connected to a functional network backbone. Yet because the two chassis are operating as a single unit, the configuration, troubleshooting and maintenance of the dual-chassis core is no more complicated than would be the case for a single chassis.

VCStack™ VCStack™ LD VCStack PLUS™

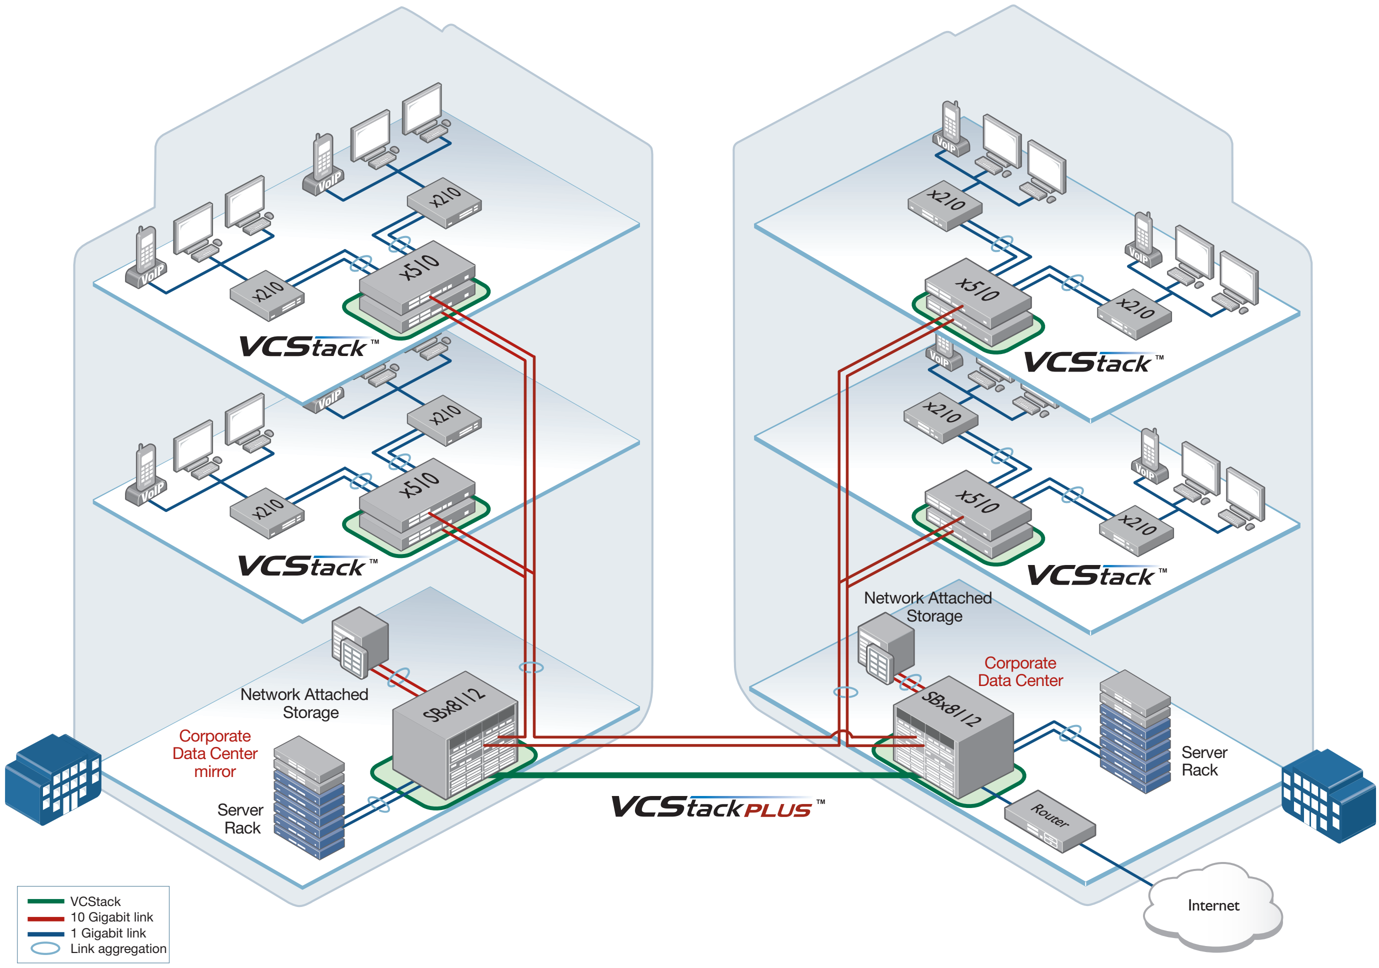

Distributed collapsed backbone

Summary

Allied Telesis provide powerful solutions for business networks. A cohesive product portfolio can support distributed or collapsed backbone environments. With our distributed collapsed backbone solution, based on VCStack technology, the best of both worlds is now available and can be implemented to meet the needs of any size business.

VCStack technology is built right in to our comprehensive feature-rich AlliedWare Plus operation system.

For more information visit www.alliedtelesis.com/alliedwareplus

About Allied Telesis

For more than 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

Allied Telesis™

C613-08022-00 RevB