Solutions Guide

Allied Telesis™

Building distributed network cores with Allied Telesis switches

the **solution** : the **network**

Introduction

Modern networks have become a fundamental component of the way we interact and conduct business. A highly reliable and high performing network is now essential in most activities, from education to data management to commerce.

Today's students make use of technology like never before, and access to online resources and applications is seen as an integral part of the learning experience. Lectures on demand, online e-learning activities and access to the Internet are expected norms in a campus environment.

Businesses also rely on online resources more than ever, with everything from standard access to email and servers, to business critical applications requiring a high availability network. This is even more pronounced where data storage is employed. Access to important information must be continuous. Multiple copies of critical data are often kept to guarantee availability and minimise any possibility of loss.

Allied Telesis provides advanced distributed network core solutions, which ensure data availability in campus environments, and also for businesses utilising data and service replication between separate sites.

Distributed network cores

The decision to implement the network core in a physically distributed form can be driven by simple considerations of feasibility, and/or by the requirements of disaster recovery.

In a campus environment, it is frequently not feasible for the network core to be in a single physical location. Departments as diverse as engineering, graphic design, science and administration all require large volume online storage and applications. Separate localised servers are often employed to meet specialized needs, and differing service requirements. Therefore, it is often natural for a campus network to consist of a distributed set of network cores, with high-speed, resilient connectivity between the individual sub-cores. This is shown in the diagram below.

The requirements of disaster recovery lead to network cores, and even network operation centres, being duplicated at 2 or more physically separated locations. This provides consistent availability of information and the knowledge that critical data is always right there when it's required. Complete data and service replication between the separate sites enable operations to continue even if one (or more) of the sites is completely out of service. The speed and reliability of the data replication between the sites is critical to the success of a disaster recovery solution. If there are significant discrepancies between the data sets at the different sites, then failover will not be seamless.

Distributed network core solutions require high performance and high availability, without adding unnecessary layers of complexity to the network or its management. The ideal network for a distributed core environment will provide additional reliability and performance, while maintaining ease of use.

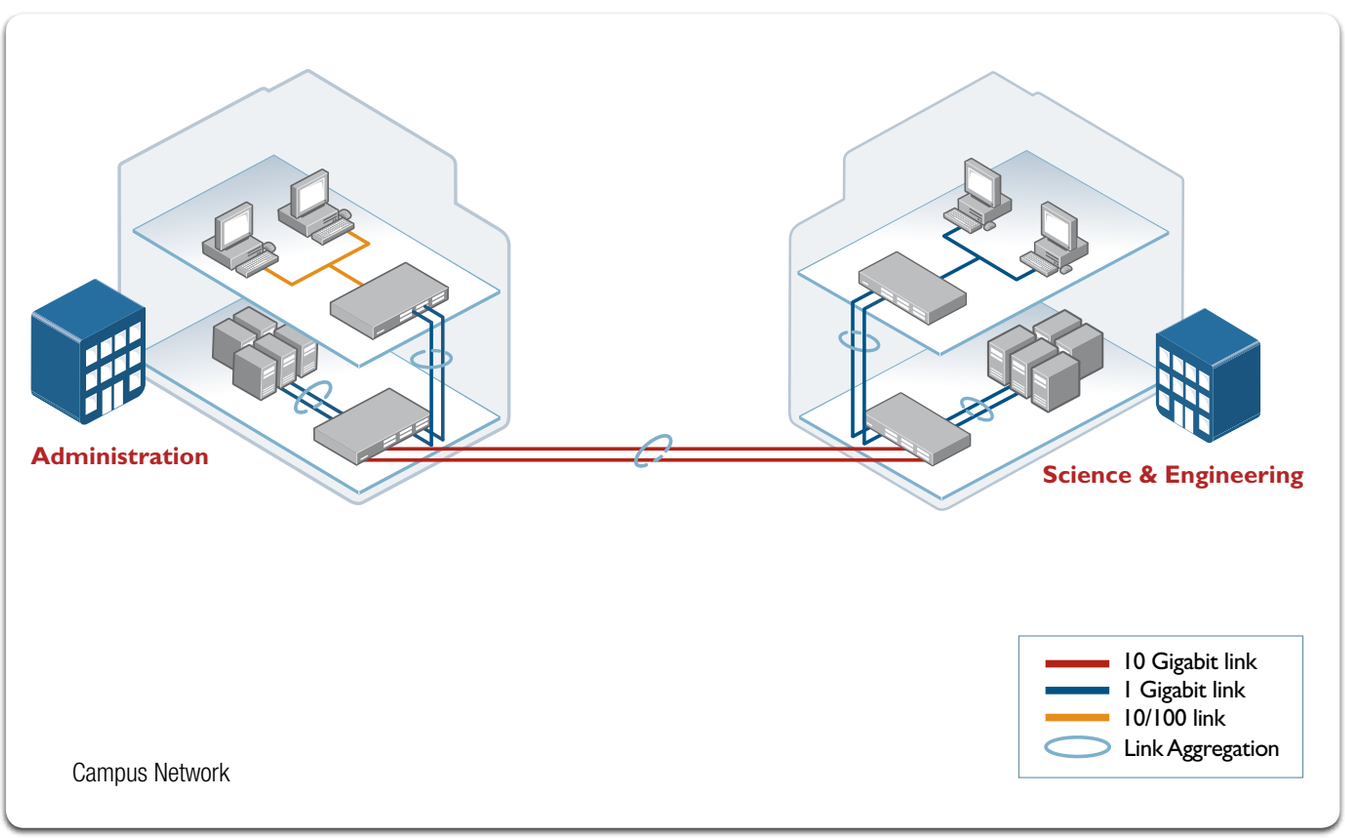

Campus Network

Allied Telesis distributed core solutions

Allied Telesis provides two technologies upon which a distributed network core can be built:

- EPSR
- Long Distance Stacking

Let us first take a brief overview of each of these technologies, and then move on to consideration of the benefits that each technology provides.

EPSR

Ethernet Protection Switching Ring (EPSR) is a mechanism for preventing data loops in ring topologies. It has a very fast recovery time in the event of broken links or switches going out of service. The recovery time can be as low as 50ms.

EPSR can be implemented on any Ethernet ring – fiber or copper – operating at any data rate. It is fully interoperable between any Allied Telesis x-series switches, iMAP products and x3100 series chassis.

The network topology within which EPSR can operate is not confined to a single ring. The Super-Loop Protection feature, an extension of EPSR, enables it to operate across a more complex topology. It can operate on multiple rings, with multiple connection points between any pair of rings. This provides a very flexible solution that can be tailored to suit any scenario. High performance is available in bandwidths right up to 10s of gigabits per second.

The diagram below shows EPSR used in a campus environment, with multiple rings providing high-speed connectivity between the various departments.

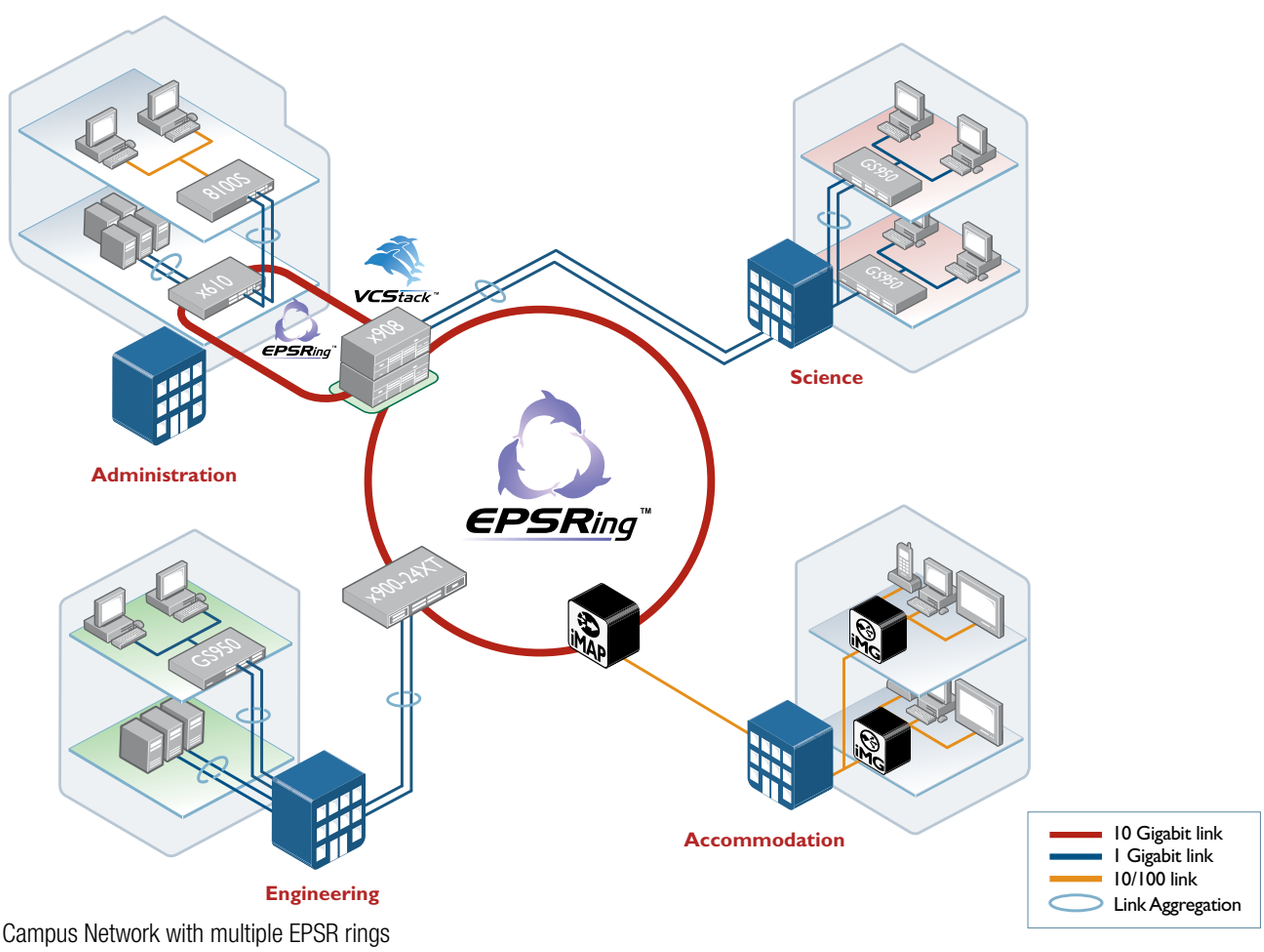

Campus Network with multiple EPSR rings

Long Distance Stacking

With Virtual Chassis Stacking (VCStack), Allied Telesis provides a truly resilient network. In normal operation, all bandwidth and all routing power in the network are fully available for use all the time. If a link or device fails, some of the bandwidth or forwarding power will be lost, but the network will still be fully operational and all remaining resources will continue to be fully utilized.

Long distance stacking enables the VCStack solution to provide a distributed network core. The increased distance provided by fiber stacking connectivity means that members of the virtual chassis do not have to be collocated. Instead, they can be kilometres apart. The distance is limited only by the capabilities of the SFP+ devices that are driving the signal down the fibre cables that connect the switches.

All of the benefits and powerful features of VCStack remain exactly the same. So the switches in a long-distance stacking solution all operate as a single switch – with completely shared software and hardware forwarding tables, a single shared configuration script, support for link aggregation and port mirroring between stack members and so on.

The diagram below shows Long Distance Stacking used to connect 2 mirrored data centers together. This single virtual distributed core ensures high availability of data for network users.

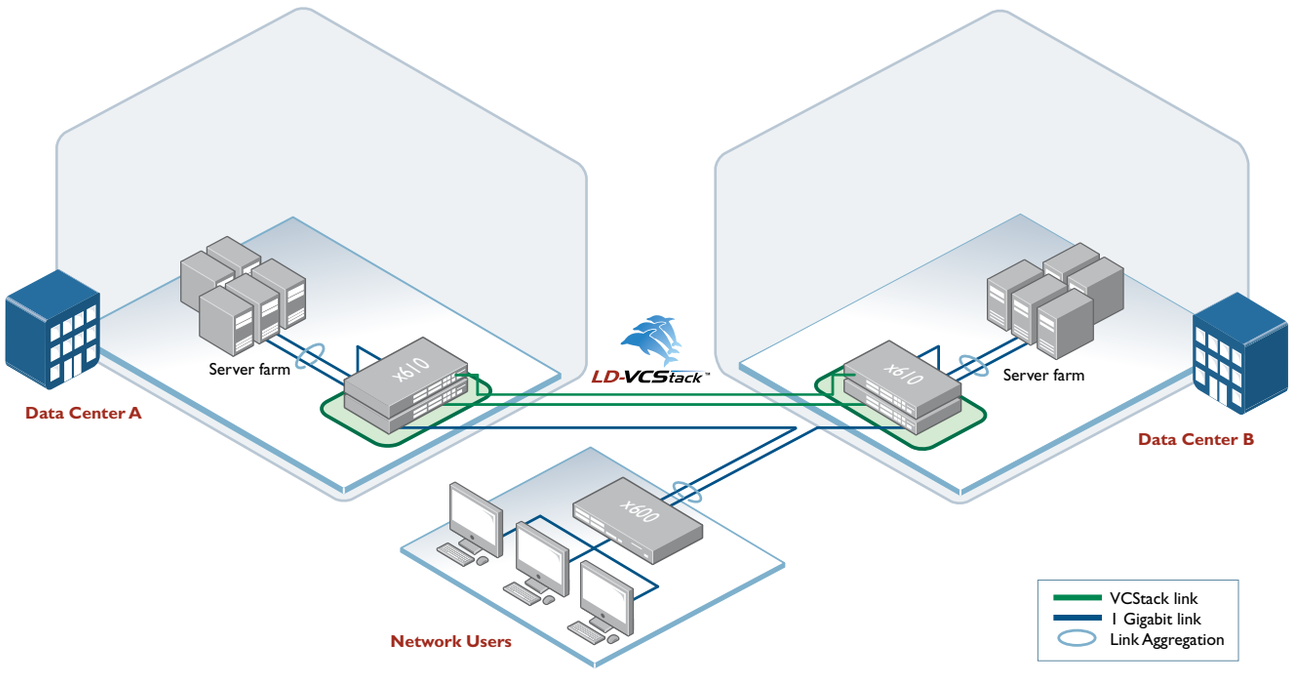

Data mirroring solution using Long Distance Stacking

Allied Telesis distributed core benefits

EPSR and Long Distance VCStack are both excellent options for building a distributed network core. They do, however, provide different benefits in different scenarios.

The prime benefits of Long Distance Stacking

Near zero traffic loss on link or unit failure

If access/distribution switches are connected to multiple stack members by resilient links, then the time required for traffic to recover from one of those stack members going down is simply the time to failover over from using both resilient links to using just one. This can be almost instantaneous. The diagram shows link aggregation between the core stack members and access switches. This ensures no perceptible disruption in the case of a link or device failure. Link aggregation is also used to connect network resources, such as servers, across the virtual chassis members.

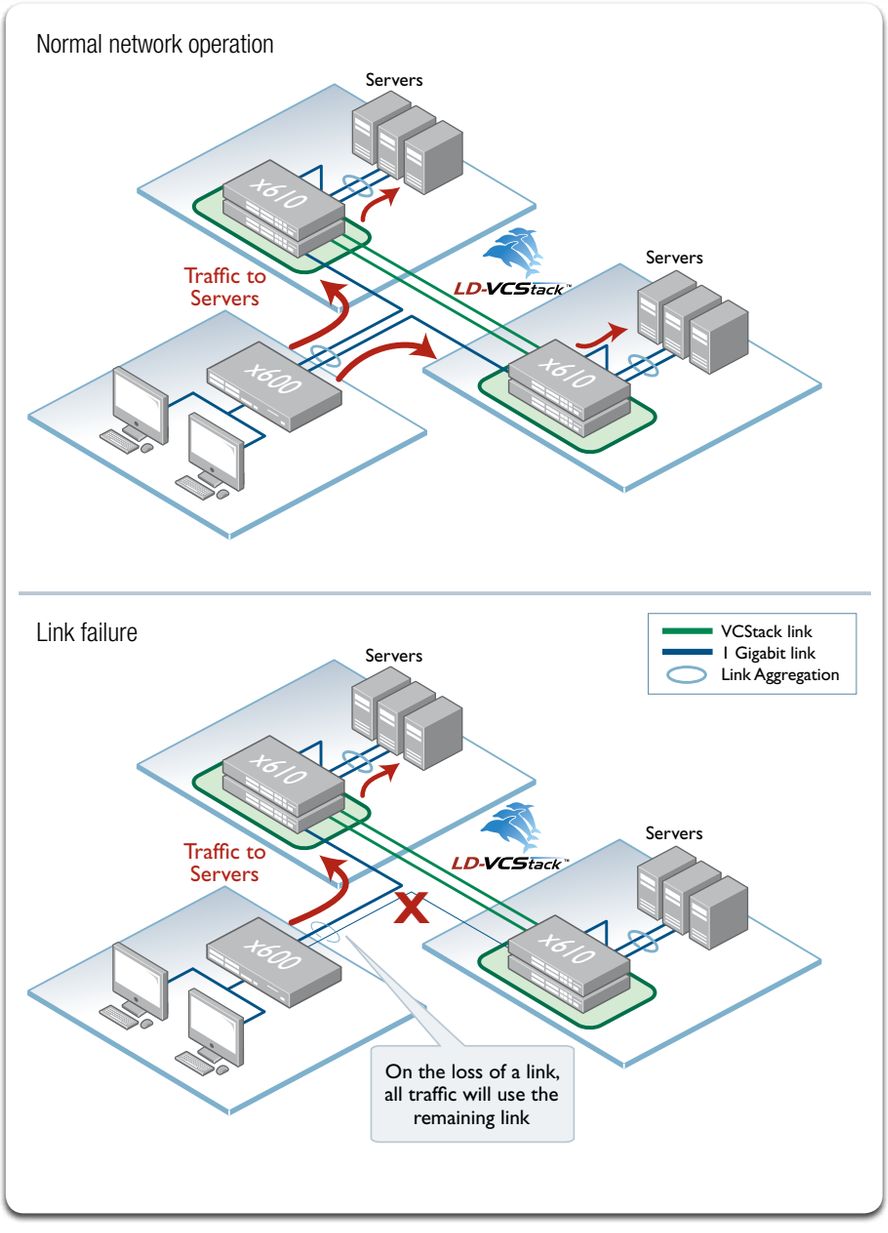

Benefits of Long Distance Stacking

Ease of management

The core long distance VCStack operates as a single switch, and is managed as a single device. Logging into any stack member gives instant access to the management interface of all the stack members. The forwarding tables, port states, and other operational information of the stack members are all presented together, as the stack is operating as a virtual chassis.

Pre-configuration of network devices

To add flexibility to the management of a VCStack, provisioning provides the ability to pre-configure (or configure 'offline') the switch ports of devices that are not currently physically present. This allows a network administrator to configure the ports of an additional VCStack member before it is actually hot-swapped in. On the physical addition of the unit, the configuration is automatically applied. This minimizes network disruption.

Simple stack member replacement

If a stack member has to be swapped out, then the replacement unit can simply be connected in its place, and will automatically join the stack, and receive the full stack configuration script. This, of course, includes the portion of configuration that is relevant to that unit, which it will automatically install and run. This facilitates effortless hot-swap of units if required.

Protocol simplification

No layer-3 unicast or multicast routing protocols need to be in operation between the core switches. The forwarding table synchronization that is inherent within the stack formation means that no routing protocols need to be configured for advertising route tables between the core switches. A VCStack core acts and responds as a truly virtual chassis.

One limitation of the Long Distance Stack solution is the fact that the bandwidth of the stacking links is limited to 10gigabits, and cannot be expanded. If all data that originates in one site, and needs to be transmitted to another site, is transported over the stacking link, that puts a limit on the inter-site traffic. However, the provision of resilient links from access switches to multiple stack members can also be used to take traffic load off the stacking links. This method of bandwidth management comes at almost no penalty in terms of failover time.

Benefits of EPSR

The prime benefits of EPSR are:

Support for a mixed-product environment

The switches comprising the distributed core can be any combination of x-series switches, x3100 chassis and even iMAPs. This is rather more flexible than Long Distance Stacking, which requires all the switches in the core to be x610-series units. In particular, if high port density is required at each location on the ring, then SwitchBlade x908 stacks (each providing over 190 ports) can be used in an EPSR ring. Allied Telesis iMAP models provide even greater port density, and also exceptional versatility, with the ability to aggregate different connectivity media in a single chassis, such as Ethernet alongside ADSL.

Flexibility in ring topology and size

There is no limit on the number of units that can be connected in an EPSR ring, whereas Long Distance Stacking is limited to a total of 8 units. Moreover, EPSR can be used in a multiple-ring topology, giving a higher level of path resiliency. This allows for large distributed networks to be created, meeting the needs of business sectors such as transportation.

High bandwidth between the core switches

When using x-series switches in an EPSR ring, the links within the ring can be aggregates of 10-gigabit links. There is no specific limit on the number of links that can be aggregated into a single EPSR ring link. So when using SwitchBlade x908 switches, with multiple 2x10gig modules, the potential bandwidth in the ring links is several 10s of gigabits.

Note that an EPSR solution does not provide the Virtual-Chassis benefits that Long Distance Stacking provides. Each unit in the EPSR ring is a separately managed unit, and there is no synchronization of forwarding tables or configurations between the units.

EPSRing™

Summary

A distributed network core is an ideal method for supporting disaster recovery solutions and large campus environments. Allied Telesis provides two leading technologies for building highly reliable distributed network cores - Long Distance Stacking and EPSR.

Each of these technologies has its own advantages in terms of ease of management, configuration simplicity, scalability and recovery time. Depending on network requirements, one or other of these technologies will provide an ideal solution for a wide range of distributed-core networks.

BENEFIT	LONG DISTANCE STACKING	EPSR
Near zero traffic loss on link or unit failure	✔	✔
Ease of management	✔	
Pre-configuration of network devices	✔	
Simple replacement of failed units	✔	
Protocol simplification	✔	
Support for a mixed-product environment		✔
Flexibility in ring topology and size		✔
High bandwidth between core switches		✔

Providing exceptionally high network availability and simplicity of operation, Allied Telesis products and advanced features have the ability to guarantee access to information when it's needed.

Products

The following products all support EPSR. The x610 series supports Long Distance Stacking.

SwitchBlade® x908

ADVANCED LAYER 3 MODULAR SWITCHES

The Allied Telesis SwitchBlade™ x908 advanced Layer 3 modular switch offers high flexibility and port density in a small physical size, providing scalable and versatile switching solutions for today's Enterprise networks. Each chassis supports eight high-speed 60Gbps expansion bays, and can also be paired in a VCStack.

x900 Series

ADVANCED LAYER 3 SWITCHES

The Allied Telesis x900 series is one of our most advanced series of switches and is unmatched in performance, flexibility and reliability. This series provides fine service provisioning granularity, high availability, and advanced QoS. All x900 switches incorporate a switching core that yields wire speed IPv4 and IPv6 routing and most switches feature robust hardware with dual hot swappable power supplies. The x900 series switches share the same expansion modules (XEMs) as the SwitchBlade x908.

x600 Series

ADVANCED LAYER 3 GIGABIT STACKABLE EDGE SWITCHES

The Allied Telesis x600 series is an advanced series of stackable switches providing high performance, flexibility and reliability. The x600 series provides high levels of Network Access Control (NAC) security making them ideal for the entry point into a corporate network. High speed 10Gbps uplinks and high speed stacking of up to four switches ensures excellent performance, while optional redundant power supplies can ensure network availability.

x610 Series

ADVANCED LAYER 3 GIGABIT ETHERNET STACKABLE SWITCHES

The Allied Telesis x610 series is the high performing and scalable solution for today's networks, providing an extensive range of port-density and uplink-connectivity options. With a choice of 24-port and 48-port versions with optional 10 Gigabit uplinks and PoE+ ports, plus the ability to stack up to eight units, the x610 Series can connect anything from a small workgroup right up to a large business.

SwitchBlade® x3112

ACCESS EDGE CHASSSIS SWITCH

The Allied Telesis SwitchBlade x3112 is a 12 Slot access edge chassis switch primarily targeted for service provider fiber access networks, and equally at home at the enterprise network edge and the data center. The switch was designed to deliver high availability, maximum performance with wire speed non-blocking backplane performance, and high port count.

iMAP 9000 series

INTEGRATED MULTISERVICE ACCESS PLATFORM (IMAP™)

The Allied Telesis iMAP is the only carrier-grade, IP/Ethernet networking solution in the industry that effectively supports IP Triple Play services simultaneously over copper and fiber infrastructure. Deploy high-bandwidth IP Video, POTS, VoIP and data services over copper or fiber using a variety of access technologies. Allied Telesis iMAPs are available in different sizes and densities and offer many plug-and-play access capabilities.

About Allied Telesis Inc.

Allied Telesis is a world class leader in delivering IP/Ethernet network solutions to the global market place. We create innovative, standards-based IP networks that seamlessly connect you with voice, video and data services.

Enterprise customers can build complete end-to-end networking solutions through a single vendor, with core to edge technologies ranging from powerful 10 Gigabit Layer 3 switches right through to media converters.

Allied Telesis also offer a wide range of access, aggregation and backbone solutions for Service Providers. Our products range from industry leading media gateways which allow voice, video and data services to be delivered to the home and business, right through to high-end chassis-based platforms providing significant network infrastructure.

Allied Telesis' flexible service and support programs are tailored to meet a wide range of needs, and are designed to protect your Allied Telesis investment well into the future.

For further information visit us online at **alliedtelesis**.com

Allied Telesis™ the **solution :** the **network**

Americas Headquarters | 19800 North Creek Parkway | Suite 100 | Bothell | WA 98011 | USA | T: +1 800 424 4284 | F: +1 425 481 3895
Asia-Pacific Headquarters | 11 Tai Seng Link | Singapore | 534182 | T: +65 6383 3832 | F: +65 6383 3830
EMEA Headquarters | Via Motta 24 | 6830 Chiasso | Switzerland | T: +41 91 69769.00 | F: +41 91 69769.11

alliedtelesis.com

C618-31026-00 RevA