White Paper

Allied Telesis AlliedWare Plus™ Operating System

Introduction

AlliedWare Plus is the latest operating system for Allied Telesis high end switch and router products. AlliedWare Plus has been in commercial use since 2008. A replacement for previous generations of operating systems, AlliedWare Plus was specified and built to be a system with the best possible characteristics to meet the challenges of future networking requirements.

AlliedWare Plus can be deployed on a wide variety of network devices, both large and small in terms of memory and CPU resource. This flexibility is achieved as much by design and architecture, as it is by being built from components that themselves support a wide variety of platforms.

AlliedWare Plus can take on new code from a wide variety of sources, including Commercial Off-The-Shelf (COTS) software, open-source projects, and code written within Allied Telesis. While the modular architecture of AlliedWare Plus is important in reducing the effects of all this code on other parts of the code base, it is still a complex undertaking that requires best-of-breed development processes and tools, as well as extensive testing. The development of AlliedWare Plus is mostly carried out using Scrum, which is a highly productive and reliable Agile software development framework.

This white paper describes the characteristics and features of AlliedWare Plus that meet the modern challenges for networking devices, and also explains how and why AlliedWare Plus is so well-positioned for future developments. It provides ample evidence of the effectiveness of the AlliedWare Plus operating system, both right now, and well into the future.

Kernel

AlliedWare Plus is fundamentally based on the Linux kernel. This is responsible for basic process control in the system, which includes starting and ending processes, scheduling process execution and allowing processes to communicate with each other. The kernel is also responsible for the management of other resources in the system, such as memory and access to hardware devices.

The Linux kernel is an example of a 'monolithic kernel', as opposed to a 'microkernel'. A lot of functionality resides in the kernel, so the ability to build on the kernel's capabilities simplifies the task of developing value-added features in AlliedWare Plus.

Advantages of Linux

Using Linux as the kernel for AlliedWare Plus has many advantages, most of which are a direct result of its popularity. Advantages include:

- Linux has a large market share in the embedded and data center markets. These are markets that AlliedWare Plus is operating in, so the use of Linux is a good fit.
- As an open source project, Linux has a large, passionate community supporting it, while retaining a centralized control model to keep quality levels high.
- The open source nature of Linux makes it a popular platform for experimental development. A number of new operating system features are first appearing on Linux platforms, and in many cases are not supported at all on other systems. The ability to pick up a new feature and rapidly assimilate it is a great advantage.
- Linux is also well supported by commercial chip manufacturers. Device drivers for the chips used in Allied Telesis products are often provided only for Linux platforms. Easy access and assimilation of these drivers means that Allied Telesis can focus instead on designing the networking features that provide customer value.
- Linux is very portable. It operates on a wide range of devices, from small embedded controllers to supercomputers. This makes for a very reliable system, and one which is continuously being tested on millions of computer systems around the world.
- The popularity of Linux means that it is familiar to the majority of the user community. The internal data tables that Linux uses, and the logging messages it produces, are well known to Linux users. Monitoring switch operations, and understanding event reports, are simple tasks for the majority of network engineers. Moreover, there are considerable resources available on the Internet to assist in understanding the logging and statistics generated by the kernel.

In summary, Linux is an excellent choice for AlliedWare Plus. Allied Telesis can build on a stable, fully featured system, in order to deliver much better value for customers with AlliedWare Plus.

Architecture

The architecture of any computer system is focused less on what the system code does, and more on how it does it. How is the code organized? What underlying mechanisms does it use to carry out fundamental tasks? How well does it perform? Although large parts of AlliedWare Plus are based on the Linux operating system, there are still significant aspects of the architecture that must be executed correctly in order to make AlliedWare Plus as good as it can be. These are explained in the following 3 sections.

Modular architecture

Most importantly, the architecture is modular. Each different feature operates in its own process in the system, controlled and protected by the kernel and using the kernel to communicate with other processes. If a process is correctly checking its inputs, it cannot be disrupted by the failure of any other process.

In systems without modular architecture, an error can result in the entire system crashing. AlliedWare Plus has modular architecture, so the effects of an unhandled software event are confined to the process within which the event occurred. Other processes continue to run and the system functions as before. The affected process can be restarted and returned to the running system.

Another major benefit of modular architecture is that individual processes can be reinitialized without bringing down the entire system. In AlliedWare Plus it is possible, under manager control, to restart the Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) processes without other disruption in the system. In fact, routes learned through these processes remain in the device's forwarding tables while the process is being restarted, ensuring no disruption to traffic at all.

As a result, AlliedWare Plus supports the Graceful Restart (also known as non-stop forwarding) extensions that are defined for these routing protocols.

Hardware abstraction

Hardware abstraction is another significant aspect of the architecture of embedded systems. A modular architecture is a necessary prerequisite for effective hardware abstraction, but it is not the whole picture. There are also the considerations of how to model different parts of the system in different ways, and how to translate between models.

Some parts of AlliedWare Plus concern themselves with the routing protocols, for example Protocol Independent Multicast (PIM), OSPF and BGP. Their operation is expressed in terms of concepts such as routes, next hops and interfaces. These are high-level representations of what lower-level software modules will write into hardware to set up forwarding paths. Hardware abstraction means that different hardware types can be supported without having to change the routing protocol code, or the abstractions that they use. It is the Hardware Abstraction Layer (HAL) that converts requests expressed in these abstract terms to specific hardware related terms relevant to the hardware-level software modules.

With hardware abstraction, it is possible to support new types of hardware without changing the code of the routing protocols. In many cases the HAL is actually provided by the hardware chip vendor, making new hardware support even easier.

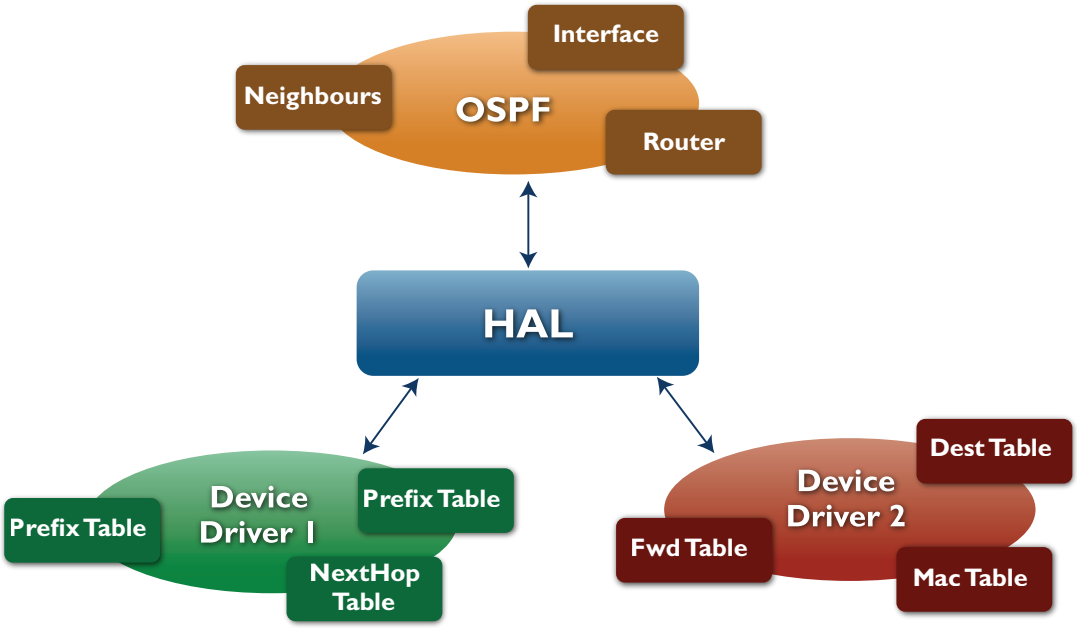

AlliedWare Plus with hardware abstraction is even more flexible. It is capable of providing more and more functionality for Allied Telesis customers, and a greater variety of platforms.

Control/data separation

A very similar concept to hardware abstraction and virtualization is "control/data separation". By using a HAL, AlliedWare Plus can totally separate the two processes that communicate using the HAL.

One example of this in AlliedWare Plus is the Virtual Chassis Stacking (VCStack™) feature. Up to 8 AlliedWare Plus switches are connected with stacking interfaces, and after election of a single master device, the control and data elements of the system are separated. The control software runs on the master, while the data software runs on each of the stack members. Rather than using an inter-process communication method suitable for processes on a single device, VCStack uses network protocols to communicate from the stack master to the other stack members.

The control/data separation implementation in VCStack is proven in the field, having been actively deployed for several years now. It has proven ability to control the full spectrum of Layer 2 and Layer 3 data forwarding, including multicast and IPv6.

As a result, AlliedWare Plus is naturally well-positioned for the development of Software Defined Networking (SDN). The AlliedWare Plus control plane is inherently able to operate independently of the data plane. The implementation of the separated control and data planes is robust and field-hardened, and proven to scale up to networks of thousands of users.

Scalability

A network consists of multiple layers of equipment operating in unison to carry data from one endpoint to another. The equipment operating in one network layer can differ greatly from the equipment operating in the next layer. Towards the core of the network, the price, complexity, forwarding performance and port count of the equipment increases, often quite sharply.

Irrespective of the difference of equipment scale at each layer, a true network solution provides a consistent user experience and feature compatibility - right across the network.

To this end, Allied Telesis provide the same AlliedWare Plus operating system across a full range of network equipment, from small 10-port Layer 2+ edge switches, up to a 12-slot chassis that can support tens of 10Gbps ports.

Evolving an operating system to operate across such a wide range of equipment is a challenging process.

To operate on low-end equipment

The software has been systematically optimized:

- data structures are optimized to achieve efficient memory usage
- replication of data between different modules has been reduced

The Inter-Process Communication (IPC) between different modules has been extensively analyzed and reworked, to fit with the principle that the level of IPC activity should only increase linearly with the number of interacting modules, rather than increase as the square of the number of interacting modules.

Boundary conditions have been extensively probed to eliminate the issues that can occur on resource-constrained platforms.

Not only has this work resulted in consistently reliable operation of the very large installed base of low-end equipment running AlliedWare Plus, it has also contributed significantly to the overall robustness and efficiency of AlliedWare Plus.

To operate on high-end equipment

At the high end of the equipment scale, the challenge has been to ensure the software can manage the multi-protocol complexity, and rapid, large, change-overs that characterize the core of a large network.

The Layer 3 switch at the hub of a large network must be utterly reliable at maintaining control of forwarding and routing tables containing thousands of entries. These entries are often generated by multiple different unicast and multicast protocols, and can have complex dependencies on each other. A state change on a core network link can require that these tables are rapidly updated as routes are learnt or withdrawn, or redistributed between protocols, or as the forwarding trees for hundreds of multicast streams require adjustment. These updates must be performed rapidly and in a completely error-free fashion.

Scaling AlliedWare Plus up to meet these core-network requirements has been made possible by a mixture of strong modular design, and multiple iterations of systematic stress testing to hone that design to the necessary level of high reliability.

Portability

To take advantage of different processors that are suitable for different network applications, an operating system must be portable across processing platforms. AlliedWare Plus is eminently portable across processors. This is partly due to the inherent portability of the Linux kernel, and partly due to the architecture and coding discipline that has been applied in the development of AlliedWare Plus.

The strong hardware abstraction within the AlliedWare Plus architecture ensures that processor-specific software components are strictly contained and clearly identified. As a result, the demarcations between processor-related software components and processor-agnostic components are unambiguous.

Within the code, "endian agnosticism" is strictly adhered to, enabling straightforward porting between processors of different endian design. AlliedWare Plus runs reliably in a mixed endian environment, where processors of different endian design are operating together in the same system.

As multicore processors are being increasing utilized in embedded systems, AlliedWare Plus has moved smoothly into the multi-core environment. The multi-threaded nature of the AlliedWare Plus software design, along with the reliable mechanisms within the Linux kernel for allocating different threads to different cores, have enabled AlliedWare Plus to unlock the performance improvements that come with true multicore processing. As the core counts of processors increase, new processor-intensive features will be added to AlliedWare Plus. These can be given their own dedicated cores, and will not impact the performance of existing time-critical functionality.

The effect of this portability is that Allied Telesis can choose the optimum processor for each hardware platform – achieving excellent performance/cost ratios, quickly taking advantage of newly available processor capabilities, and utilizing multi-core efficiencies – without adverse impact on time-to-market.

Development

The development processes used by Allied Telesis greatly contribute to the strength of AlliedWare Plus, namely reliability, process, and testing.

Reliability

The best way to create a reliable system is to have best practices in the development process, and the best code.

As the saying goes, "the best code is the code you don't write", and this is certainly true of AlliedWare Plus, since a large amount of the code in AlliedWare Plus comes from the Linux kernel.

Allied Telesis has always been dedicated to continuous improvement of the development process. Development teams are evolved with care to ensure that capability is retained, and are expanded in a sustainable fashion. Hiring policies create a mix of experience and state of the art knowledge. Active research is put into discovering best practices for software development. The result is the creation of an efficient, resilient, mature development system that can develop new products and features in good time, and fully support them for at least a decade. This provides customers with the assurance of consistent high quality, and the confidence that their investment will be protected for a worthwhile duration.

Process

Allied Telesis uses the very latest Agile scrum methodology to develop AlliedWare Plus. The modular architecture allows a number of teams to work independently of each other on different parts of the code. The emphasis of scrum on delivering working code means that integration of new code happens regularly and frequently, and therefore AlliedWare Plus is never far from a releasable state.

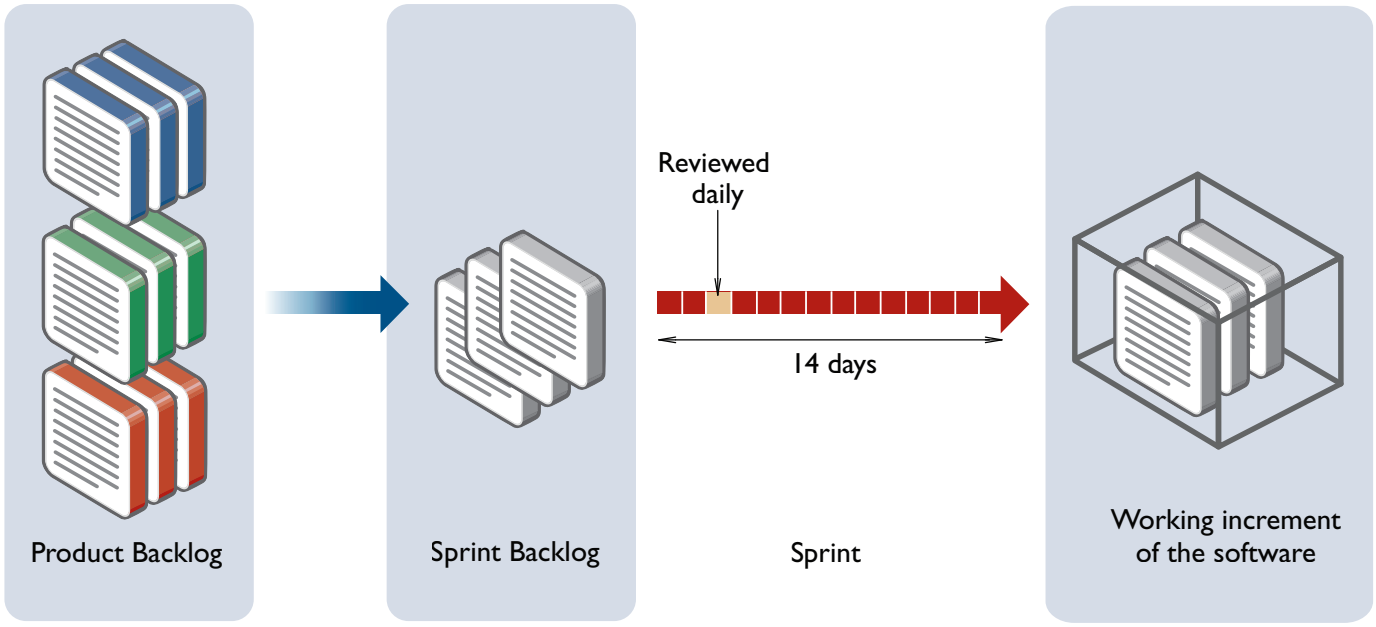

Software development teams create a product backlog, encompassing work required for delivery of a fully functional product or software feature. By working in short 'sprints', teams deliver essential features first, and are flexible enough to change priorities as and when required. Development is test-driven to create high reliability in the code and ensure that projects don't drag on, causing schedule slip and inefficiency.

Parallel development

The disciplined approach to software development has made it possible for Alliedware Plus to be developed on multiple product lines, in multiple development centres, simultaneously.

Geographically separated teams, with expertise in different market sectors, are able to develop market-specific features for their product lines using the shared Alliedware Plus codebase. This enables efficient time-to-market whilst maintaining a uniform user experience over a wide range of Allied Telesis products.

Testing

AlliedWare Plus testing occurs in multiple phases.

Developer testing

The Agile methodology embeds testers into software teams to ensure robust developer testing at the same time that the software is being written.

Continuous integration testing

Allied Telesis have adopted the model of continuous integration, which has proven powerful in the competitive environment of the modern software industry. Software developed by separate teams is integrated into the common existing code base daily. Integration testing is continuously verifying that new deliveries have not caused degradations.

Integration testing is a combination of automated tests and manual testing. Automated testing provides wide coverage of base functionality in a short time. Manual testing finds issues caused by unusual uses of the system. Rather than settling for the predictability of running the same automated tests over and over, Allied Telesis uses the creativity of its testing teams to explore outliers in the way that systems are connected, configured and stressed.

Simulated real-world testing

Software then passes through multiple phases of testing in network scenarios, both large scenarios simulating real-world environments, and more focussed scenarios designed to stress critical aspects of the software to a degree well beyond that normally experienced in production networks.

Customer feedback

Feedback from customer issues provides further product strengthening. Reviews of all customer issues are used to refine and expand the testing coverage.

"Defense in depth"

AlliedWare Plus is a feature-rich operating system. Besides the expected Layer 2 and Layer 3 network protocols, a number of features in particular make AlliedWare Plus a highly reliable and maintainable system.

Many factors work together to aid reliability and reduce downtime. By applying continuous feedback from support staff and customers Allied Telesis has defined a number of features that achieve this goal.

"Defense in depth" is a concept in which multiple layers of reliability controls are placed throughout a system, to provide highly resilient handling of expected and unexpected events. A well-designed system needs to deal with events that may arise due to software errors. Even if a process doesn't crash, it is good practice to monitor it for less serious issues that could cause problems, and restart the process if it becomes necessary. Even in normal operation, monitoring and logging unusual events benefits network troubleshooting.

Defense in depth is the rational behind the following features:

Recovering from errors

In any system, there are s number of possible responses to a module encountering an unrecoverable error. The most basic response is to do nothing, and therefore require a manual intervention to restart the affected module. The option used in AlliedWare Plus is more robust than this. First, the stack dump generated as a result of the error is saved. This can prove vital for later debugging efforts. Ample space is provided in Allied Telesis devices for storage of core dumps and other diagnostic and management files. Once the file is saved, the module is restarted.

In this way, AlliedWare Plus achieves the ideal combination of maintaining up time and gathering the required information for debugging.

Reloading modules if they have used up too much memory

There is a class of software error known as memory leaks. It is often the case that a process that leaks memory is still fully functional, but the operating system will eventually start running out of memory for other modules. Ideally, remedial action is taken against processes that have leaks.

AlliedWare Plus has a simple mechanism to check certain processes at regular intervals for leaks. A regular task checks the amount of memory used by the monitored processes. If any process is discovered to have leaked more than an acceptable amount of memory, that process is restarted. All leaked memory is recovered when the process is stopped, and the system returned to normal operation.

Process health monitoring

AlliedWare Plus also monitors the health of selected processes. A regular message is sent to monitored processes, and if not enough responses are seen, the process is marked as unresponsive.

Hardware health monitoring

Design and manufacturing processes are very strictly controlled in Allied Telesis products, but occasionally some part of the hardware can create an issue, particularly once a unit has been in operation for some years. Hardware issues are particularly serious, because they can cause very unpredictable results. Hardware health is therefore monitored very closely. A configurable range of actions can be taken by the software upon identifying a hardware problem.

Access to Linux utilities

The Linux implementation within AlliedWare Plus contains a suite of standard Linux tools. Most of the popular Linux utilities are already exposed via AlliedWare Plus Command Line Interface (CLI) commands. But, for the expert user who wants to use very specific combinations of command-line options on these utilities, then access to the underlying utilities is available via shell scripts.

Moreover, shell scripts can be created and executed from the AlliedWare Plus CLI, that can run combinations of underlying Linux commands, and also AlliedWare Plus CLI commands. This enables users to precisely tailor the set of information they gather in regular, or event-driven, monitoring activities.

Triggers and scripting

Precise network monitoring, and specialized network operations require the networking devices to provide configurable event-based actions. AlliedWare Plus provides uniquely flexible event-based configuration, in the form of triggers and scripting.

Triggers can be configured to execute on a range of event types, including:

- Peaks in CPU utilization
- Interface state changes
- VCStack master failovers
- Losing Layer 3 connectivity with a remote host

Additionally, triggers can be configured to execute at specified times. A trigger can be configured with up to 5 scripts that are activated when the trigger goes off. These scripts can be sets of AlliedWare Plus CLI commands, or shell scripts. There are absolutely no restrictions on the commands that can be included in the scripts called by a trigger.

This functionality enables you to perform extremely precise and flexible monitoring of network events. Intermittent events in a network are notoriously difficult to capture and diagnose. Being able to use the network switches as configurable probes is a significant addition to a network manager's toolset.

The ability to have switches automatically change configuration on a timed basis, or in an event-driven manner, also expands the range of business rules that can be feasibly implemented within the network.

Security and user authentication

Network security is complex. While an external Internet firewall is vital to a secure network, it is equally as vital to have a strong wall of defense against attack from within.

For this reason, AlliedWare Plus provides complete access edge security. All ports on AlliedWare Plus switches support tri-authentication: 802.1x, MAC authentication and Web authentication. As a result, all devices attached to the access edge can be authenticated, and there is no need to make exceptions for any devices or guest users. With authentication applied to every access edge port, the wired network has a complete ring of defense against unauthorized connections.

Management access to AlliedWare Plus is fully secured. All user access to the command-line management interface must be authenticated. Different levels of access are provided to different users, depending on their allocated privilege level. The software can be configured to reject any unencrypted management access.

Security does not necessitate inconvenience. In fact, some of the AlliedWare Plus security features enhance user convenience.

Roaming authentication transfers users' authentication status from port to port, if their point of access moves. For example, as a roaming wireless user moves from one access point to another, the shift is visible to the switch that the access points attach to.

AlliedWare Plus contains an embedded RADIUS server so that for small networks, the network access authentication can be contained entirely within the switching infrastructure.

Utilizing external media

Portable mass storage devices, such as USB flash drives and Secure Digital (SD) cards, provide extremely convenient ways to collect, store and transport information. AlliedWare Plus presents a number of management options to allow customers to best utilize portable mass storage devices.

The SD autoboot feature allows users to clone a switch - set up a new switch with exactly the same configuration and software version as an existing switch. New or replacement switches may be automatically configured with the required configuration and software version by inserting a pre-configured SD card and booting the switch.

Network monitoring activities can also be output to external media, which have large storage capacity to enable long periods of traffic to be captured. The switch can be directed to write tcpdump output to a USB flash drive or SD card.

Detailed security auditing can be supported by regularly dumping Dynamic Host Configuration Protocol (DHCP) snooping tables and/or 802.1x supplicant details to external media. In this manner, very detailed, time-stamped, audit data, covering long time periods, can be easily stored. Moreover, the data is stored directly at the edge switch where client devices attach, so it is not lost if connectivity to a logging server is interrupted.

Powerful diagnostics

Problems will always occur in networks. Unexpected problems can arise no matter how well a network and its equipment are designed, implemented, and tested.

The key to minimizing the disruption caused by problems is to diagnose and resolve them as quickly as possible. Devices that provide very little information about their internal state, and about the data passing through them, create barriers to efficient problem resolution. Equipment that provides software and hardware status and statistics down to a very detailed level can significantly enhance efficient problem resolution. Moreover, if that equipment supports configurable monitoring tools, then it becomes a key tool in the network troubleshooting process, and can even provide early warnings of emerging network issues, or of problems developing within its own hardware or software.

The philosophy in the development of AlliedWare Plus has consistently been to provide as much diagnostic information as possible at every level of the hardware and software. This provides a powerful monitoring tool for the end user. Additionally, it enables detailed quality assurance during the development phase - resulting in a high level of software reliability.

Industry-leading technologies

The AlliedWare Plus operating system provides a reliable and expandable platform upon which Allied Telesis have been able to build some truly game-changing networking technologies.

VCStack, VCStack Plus, and VCStack-LD

Allied Telesis is unique in providing a consistent stacking offering right from the edge to the large-network core. High port density stacks at the network edge have the same rich set of high availability features as do stacks of chassis switches at the network core.

This provides a consistent user experience across the network, and the ability to choose the most appropriate product for any given network role without being constrained by arbitrary compatibility requirements.

The modularity and expandability of AlliedWare Plus is well illustrated by the fact that such a range of high-availability features has scaled right up to stacked pairs of 12-slot dual-controller chassis. The chassis stacking, labelled VCStack Plus, achieves full table synchronization for the unicast and multicast forwarding and control planes at Layer 2 and Layer 3 for IPv4 and IPv6 across two full chassis, with up to four active controllers. This level of multiprotocol, multilayer and multiplane synchronization, coupled with quad-controller resiliency, is almost unmatched in the industry.

AlliedWare Plus deployment scalability is provided with VCStack-Long Distance (VCStack-LD), which enables stack members to be kilometers apart, supporting distributed network environments and data-mirroring solutions.

VCStack™ VCStack PLUS™ VCStack™ LD

Ethernet Protected Switching Ring (EPSRing™)

Allied Telesis provide an advanced resilient-ring solution, with failover times as low as 50ms, that operates across a large range of enterprise and metro products. This brings Telco-like network resilience into the world of Ethernet-based Metro/Campus/Enterprise networks.

By providing this capability across a wide product range, over a full spectrum of different feeds and speeds, Allied Telesis makes it possible for **any** Ethernet network to enjoy extremely rapid link failure recovery, without incurring the complexity and expense they might expect this to require.

AMF

Allied Telesis Management Framework (AMF) is a sophisticated suite of management tools that provide a simplified approach to network management. Common tasks are automated or made so simple that the every-day running of a network can be achieved without the need for highly-trained, and expensive, network engineers. Powerful features like centralized management, auto-backup, auto-upgrade, auto-provisioning and auto-recovery enable plug-and-play networking and zero-touch management.

For more information about Allied Telesis and the AlliedWare Plus operating system visit: alliedtelesis.com/alliedwareplus

About Allied Telesis, Inc.

Founded in 1987, and with offices worldwide, Allied Telesis is a leading provider of networking infrastructure and flexible, interoperable network solutions. The Company provides reliable video, voice and data network solutions to clients in multiple markets including government, healthcare, defense, education, retail, hospitality, and network service providers.

Allied Telesis is committed to 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™ the **solution** : the **network**

North America 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 & CSA Operations | Incheonweg 7 | 1437 EK Rozenburg | The Netherlands | T: +31 20 7950020 | F: +31 20 7950021

alliedtelesis.com

C613-08018-00 RevA