Service Description

Fujitsu Software Defined Networking (SD-WAN) High Assurance

G-Cloud 13

Contents

1. **Fujitsu Software Defined Networking (SD-WAN) 4**

Benefits of Fujitsu's SD-WAN Service for Service Pack 1 and 2 7

The Technical Design Explained for Service Pack 1 and 2. 7

Edge Devices 9

Service Operation 10

Support Services 10

SD-WAN Professional Services 11

2. **SD-WAN for the Law Enforcement Community, Service Pack 3 ..... 12**

2.1 Purpose of Document 12

2.2 Risks, Assumptions, Issues and Dependencies 12

2.3 Constraints (Standards, Policies, Guidelines) 12

2.4 **Core Platform** 19

2.5 Edge Devices 26

2.6 SD-WAN Overlay 28

2.7 Application Monitoring 44

2.8 Customer Portal Access 46

2.9 PSN and Internet Underlay 46

2.10 Customer LAN 48

2.11 AWS Edge 49

2.12 Availability & Resilience 51

2.13 Service Interruption 57

2.14 Backup and Recovery 58

2.15 Disaster Recovery 58

2.16 Performance Management 60

2.17 Security 61

3. **Service Delivery, Service Packs 1 & 2 ..... 67**

Service Management 67

Security and Information Assurance 67

Business Continuity 67

Training and Consultancy 67

4. **Service Levels, for Service Packs 1 & 2 ..... 67**

5. **Service Delivery, for LEC Service Pack 3 ..... 69**

Security and Information Assurance 69
Business Continuity 69
Training and Consultancy 70
Secure Destruction of Edge Equipment 70
Protective Monitoring & SOC SIEM Support 70
6. **Service Levels, LEC Service Pack 3 ..... 70**
Service Desk 71
Service Reporting 74
7. **Price Approach Service Packs 1, 2 and 3 ..... 75**
License Capacity 75
8. **Commercial Service Packs 1, 2 and 3 ..... 76**
Ordering and Invoicing Process 76
Trial Service 76
Minimum and Maximum Terms 76
Termination terms 76
Consumer responsibilities 76
Service constraints 76
Service exclusions 77

1. Fujitsu Software Defined Networking (SD-WAN)

Fujitsu's Software Defined Networking (SD-WAN), is a deployed subset of Fujitsu's wider Software Defined Networking services portfolio. Our SD-WAN offering comprises of three service pack offerings using the best of breed vendors and provides Customers with a catalogue of services and infrastructure that supports an assured and secure network overlay solution. Fujitsu's solution enables Customers to replace or to enhance the current connectivity network (WAN). The service comprises of three offerings from Fujitsu deployed on three shared platforms meeting OFFICIAL, with service and solution uplifts to support caveats (such as SENSITIVE), Critical National Infrastructure "CNI" standards and up to SECRET infrastructure deployments (note using exclusive assets). The Fujitsu SD-WAN solutions are described as Service Packs comprising of,

- Service Pack 1, a commodity offering featuring use of public cloud or vendor infrastructure with SD-WAN Orchestrators deployed in the UK and Amsterdam. Supported by Fujitsu network operations centres (OFFICIAL only).
- Service Pack 2, using Fujitsu's own dedicated infrastructure with all Orchestrators deployed in the UK in Accredited facilities. Supported by Fujitsu UK (Defence and National Security) network operation centres (OFFICIAL SENSITIVE & Above). All Fujitsu staff managing the service are Security Cleared (SC).
- Service Pack 3, using Fujitsu's dedicated infrastructure with all Orchestrators deployed in the UK in Accredited facilities for the Law Enforcement Community (LEC). Supported by Fujitsu UK (Defence and National Security) network operation centres. All Fujitsu staff managing the service are Security Cleared (SC) and Non Police Personnel Vetting (NPPV) cleared.

  (Note the service description (including diagrams) contains end user information and some contents has been "redacted". Upon identification of end user and in accordance with the Authority Security Aspect Letter, Fujitsu will provide copies of the unredacted or removed contents.)

The services provided include pre-production environments aligned to the Customer environment to support Fujitsu testing, patching and software updates.

The choice of Service Pack offering will be dependent on the security and Information Assurance requirements of the Customer. In the case of Service Pack 3 for LEC the SD-WAN platform meets the published certification requirements of PDS / NPRINT and complies with Police Assured Secure Facilities (PASF) requirements.

Prior to award of a contract specialist SD-WAN solution architects will be assigned to define the functionality and to agree with the Customer the Information Assurance requirements (suitability of Service Pack selected). This process will ensure the full scope of the service is understood which will also be summarised in a Statement of Work.

Service Pack 3 offers a mature defined service description for the Law Enforcement Community (LEC) and is included in this service definition document.

The Service Pack charges applicable for the Customer for services provided are calculated from the SFIA rate card Professional Services (design, deployment and any service management elements).

Hosting charges for Service Pack 1 is included in the user charge as detailed in the optional service catalogue.

Hosting charges will apply for Service Packs 2 and 3 as detailed in the optional service catalogue.

The Customer may elect for Service Pack 1, 2 and 3 to include the vendor SD-WAN user licences required. Or request Fujitsu operates cloud licenses from the SD-WAN vendor for the term on the Customers behalf (with all licences contracted on behalf of the Customer). This arrangement reflecting CCS arrangements with strategic vendors.

Alternatively for Service Pack 2 and 3 users the Customer may novate the benefit of current vendor licenses it holds subject to validation of term and permission of the licence owner.

For Service Pack 1, 2 and 3 the Customer may require edge device hardware deployed at sites which can be:

- Rented from Fujitsu as an optional catalogue item, or,
- Request Fujitsu manage edge devices provided by the Customer for the term.

To aid the Customer all applicable charges using the published SFIA rate card or from the optional catalogue will be summarised in the Statement of Work.

Once completed the Statement of Work shall confirm the defined services and outline deployment timescales to be performed by Fujitsu for the Customer. Upon agreement of the Statement of Work, and formal award of a call off contract a detailed implementation plan and agreed milestone dates will be provided confirming timescales and dependencies.

SD-WAN Overview, Service Pack 1 & 2

Fujitsu's SD-WAN solution allows the Customers to disaggregate traditional WAN connectivity and services or to adopt a bearer of opportunity strategy, for example supporting government Customers moving from the PSN, replacing with ISP connectivity other networks or MPLS, as well as providing secure connectivity to Public Cloud Service providers. We use the Customer provided connectivity as the SD-WAN underlay and enable the Customer to use other Public Sector Frameworks to source connectivity solutions. Fujitsu's SD-WAN service provides supplementary technology on its SD-WAN platform to help the Customer assure any connection proposed using ITHC tested edge devices (NFV), NCSC aligned encryption standards, and an ITSM tool set to comply with the relevant standards and deployed options.

The adoption of cloud-based applications in many respects exposes the limitations of traditional networks using MPLS and QoS which most government networks use today. Customers are seeking high "assured" availability but critically with real-time performance and relevant service management capabilities. The trend towards insourcing is supported by the Fujitsu SD-WAN off the shelf services and an ability for the Customer to self-manage the platform (default service) with a solution that supports application awareness capability and strict network performance criteria.

Fujitsu's SD-WAN solution is complemented by optional real-time analytic tools shared with the Customer to show true performance of the network and the applications using the service, analytic tools can be hosted in the public cloud or using the Fujitsu UK hosted solution.

To connect to the Fujitsu SD-WAN Service Packs the Customer is responsible for WAN connectivity to Fujitsu UK nominated peering points, this is usually sourced by the Customer using a Crown Commercial Service (CCS) connectivity Frameworks. The Customer will benefit from many of the current connectivity suppliers are already present at our UK peering points, enabling a cost-effective connection to be made. In addition to support transition from legacy government networks Fujitsu also has direct peering points to the PSN, which offers MPLS to SD-WAN gateways and Internet connectivity if required (this can be procured from Fujitsu using RM3808 or its successor).

Fujitsu's SD-WAN Service is a COTs approach, with interchangeable vendors and components with the functionality assured. Delivered with a choice of Fujitsu edge device hardware using a VNF configuration or the SD-WAN vendor edge devices. Selection of technology will be in consultation with the Customer but comprises of Cisco, Juniper and Fortinet. Fujitsu Edge devices if selected also can support other deployed applications (deployed as service chains) such as NGFW, additional SDN encryption capability (aligned to PRIME) and well as various test agents such as Paragon Active Assurance.

Fujitsu's SD-WAN service is proven and deployed in both assured Public and Private cloud environments and service management centres. We also can offer pre-staging procedures to support complex Information Assurance standards.

Fujitsu provides the Customer with catalogue options for:

- Locations of the SD-WAN Orchestration
- Regional ISP Breakout or Secure Internet Gateway

- High Availability deployment options
- Edge device types and secure builds
- Integral Firewalls or NGFW)
- Service Management options
- Protective Monitoring
- Security Operations Centre services

All upgrades and patching of the SD-WAN service are fulfilled in accordance with published NCSC guidelines for Critical; Important and Other updates with only approved vendor hardware and software deployed. The platform is evergreen and current (N-1).

Fujitsu with its key suppliers operates a pre-pod solution to support Customer validation testing and proof of concepts. The pre-pod solution is charged for in accordance with the SFIA rate card to test new services or functionality.

Today Fujitsu's SD-WAN Service currently supports deployed vendor solutions from Cisco and Juniper. We can provide all the licences required to support the SD-WAN functionality. Where a Customer has a Cisco transferable license or offers a license offered under Crown Commercial Service terms this can be deployed on either platform (Public or Private Datacentres) subject to validation of license terms.

To provide enhanced functionality, Fujitsu operates a R&D facility with development partners and smaller vendors that are then added to the Service once validated (Service Chain Approach).

Fujitsu's Service has benefited from over 3 years of solution development, service desk tooling and ITHC testing to support rigorous various Information Assurance requirements. Today we operate for the UK government some of the largest SDN based networks in the UK and abroad. Fujitsu SD-WAN Service and security approach fully aligns to MEF and the UK National Cyber Security Centre (NCSC) published guidance or standards.

The following catalogue or service options are available from Fujitsu, a full description for the Customer will be provided in the Statement of Works upon request.

The Fujitsu SDN SD-WAN solution supports a transparent on-boarding and definition of service, prior to award with specialist SD-WAN solution architects offering expertise with the following design attributes:

- Resilient (High Availability by design) Orchestration Platforms
- Various Orchestration hosting options (to meet Security Requirements)
- A Boundary Protection Service of infrastructure with Customer service uplifts
- Managed Service Options (ITIL aligned) from Customer self-managed, to a fully managed service provided by Fujitsu
- CNI aligned measured uptime of core platform measured over 24 hour 365 days with service desk options to support Customer requirements
- Analytic tool sets to view Customer applications in transit and network performance
- Protective Monitoring of platform and variations aligned to Customer requirements
- Enhanced SOC services
- A charging model based on edge device selected and license capacity per site
- A fully featured SD-WAN solution
- Emergency deployment solutions supporting pop up sites and mobile locations
- Supporting from 2 to 2,000 sites per instance
- Automated capacity management
- Testing of new vendor functionality (updates) before deployment supporting an evergreen solution
- Choices of vendor edge devices supporting up to 10Gbps
- NFV based edge devices supporting from 10Mbps to 10Gbps at a site

- Choice of standard or High Availability edge site deployments
- Network Encryption aligned to Foundation
- Network Encryption aligned to PRIME via NFV or edge hardware
- Alignment to Customer Information Assurance policies (including Customer ISMS)
- Design, Deployment and Project Management services using SFIA catalogue
- Transition Service and assistance at termination of contract (including use of gateways)
- Hypervisor connectivity to the main Cloud providers

Benefits of Fujitsu's SD-WAN Service for Service Pack 1 and 2

Fujitsu's SD-WAN Service encompasses many benefits including:

- A COTs approach using leading vendors technology and functionality to meet Customer requirements
- Unlocks access to 30-40% savings from removing dependency on legacy infrastructure or services and enables faster transition from legacy networks
- Supports an insourcing strategy which can be phased in as experience is gained in running an SD-WAN solution
- Meets Customer Information Assurance requirements with design, documentation and deployment skills
- Cost effective start-up and deployment
- A catalogue approach to services and functionality (selecting services required)
- Separation of Security Functions from Network Management for High Assurance deployments
- Evergreen with updates and functionality added automatically (with controls) during term of agreement
- Choice of Service management Options (Customer Managed or full Fujitsu Management)
- Access to Fujitsu assurance teams to design and operate to Customer Information Assurance policies using SFIA rate card
- Overlays onto Customer switching fabric without feature loss such as MPLS
- Fujitsu or Cloud hosting options and use of other datacentre deployments
- 24x7x365 (co-located 1st, 2nd and 3rd Line and Service Management in the UK)

The Technical Design Explained for Service Pack 1 and 2.

Fujitsu's SD-WAN service is Orchestrator based, service providing a single plane of glass for the Customer or Fujitsu for service management. Access to the features and program options are based on agreed RBAC's between the Customer and Fujitsu and will be dependent on the service management options selected (Customer Management option or Fujitsu Managed Service). The SD-WAN features and functionality will be driven by the application deployed; however, our solution approach usually enables full feature access to vendor functionality (subject to any limitations imposed by Information Assurance) the Customer once deployed. Selection of vendor technology can be by the Customer or using Fujitsu expertise.

The Orchestrator manages all connectivity circuits, Edge devices, policies and provides real-time information of the availability of the Service to the authorised user. An intuitive drag and drop design including configuration options allows the Customer to see all status in real time.

It uses configurable exception reporting indicating where performance of a bearer at a site or across a network has reduced below allowable tolerances (say for example Packet or Jitter loss), site brownouts and where security policies need to be reviewed.

Regardless of the service management arrangements (Customer or Fujitsu management) access to real time performance and status is provided to the Customer via a Customer Portal capability and/or API.

Fujitsu's SD-WAN solution, will by policy, seek to maintain service and will implement workaround solutions if components or underlining connectivity is not available or performing below predefined tolerances. Upon restoration of the Service impacting the solution, our SD-WAN service will revert to the default configuration

automatically without user intervention. Fujitsu's SD-WAN Service enables suitably qualified Customer personnel to be able to manage the Service, with policy creation and management using standard scripts.

For application aware based routing (routing on first packet), a database is fed to the Orchestrator of known applications, which is updated monthly. Depending on vendor selection, this feature will suggest to the Customer the best policies to be adopted across the known bearers. Fujitsu also can provide optional discovery tools (via catalogue) to identify other applications in use in the Customer infrastructure and to create a table of suggested priority.

Fujitsu's Orchestrator functionality includes with rollback options and the ability to schedule updates by region or time. This means that once changes are made, they are automatically uploaded to all edge devices deployed in the Customer database. The health of all devices managed by the Customer is shown including connectivity status, performance status, firmware status and certificate status. This is most useful in cloud environments when application performance can be monitored.

Fujitsu's Orchestrator design supports HA requirements offering 99% to 99.99% availability, depending on the hosted solution deployed, and the operating software deployed. Regardless -0If Public or Fujitsu Cloud solution adopted the solution enables the edge device to continue to operate (assuming connectivity is available at a site) based on its current configuration for periods defined by the Customer even if access to the SD-WAN Orchestrator is unavailable.

The following provides a typical schematic of the functional layers and service models provided.

Note: some changes may apply depending on the SD-WAN application deployed.

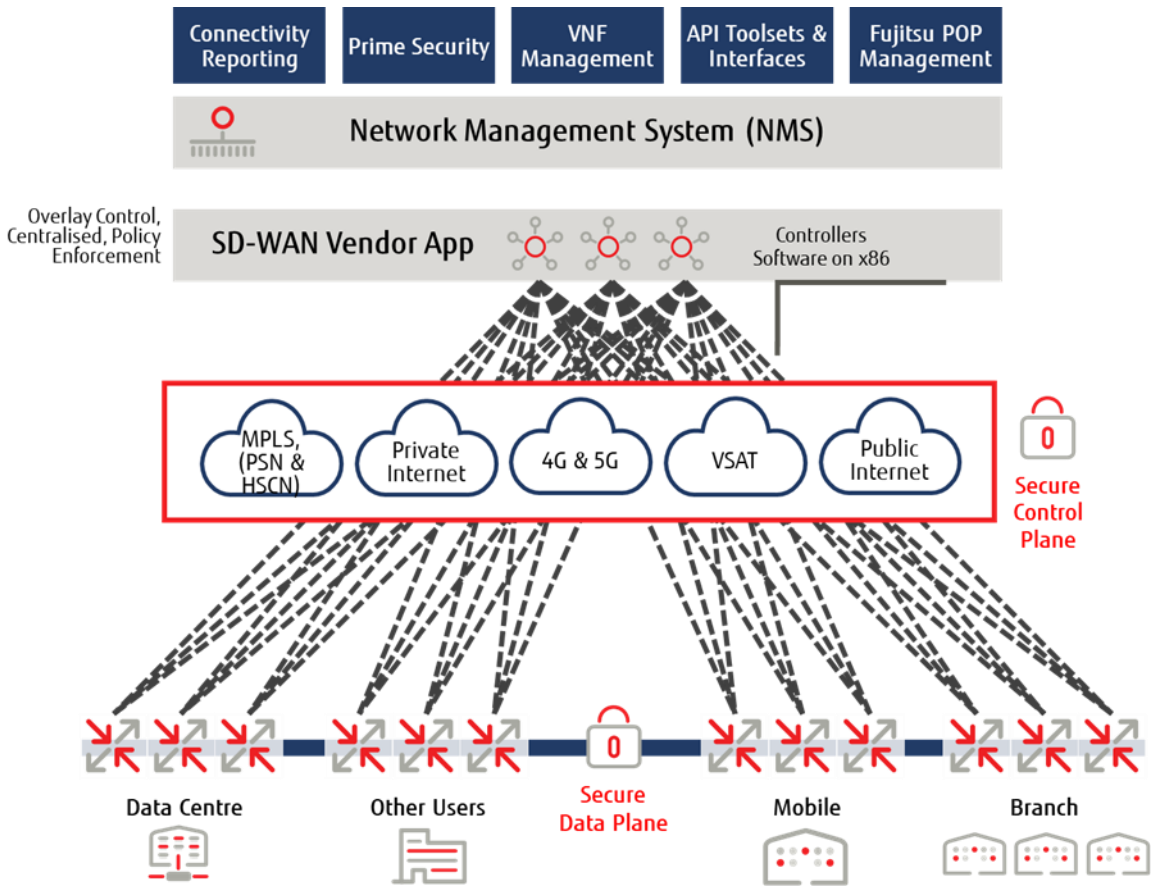

Figure 1

Edge Devices

Fujitsu's SD-WAN Services provides two edge device options:

1. **Fujitsu Edge Devices, supporting VNF based applications (standard offering)**

Fujitsu edge device consists of the following components:

- OpenStack
- NFV MANO
- KVM
- DPDK
- Optional Virtualized Network Function (supporting SD-WAN Router, NGWF, IDS/IDP etc.)
- Standard or High Availability appliance deployments.

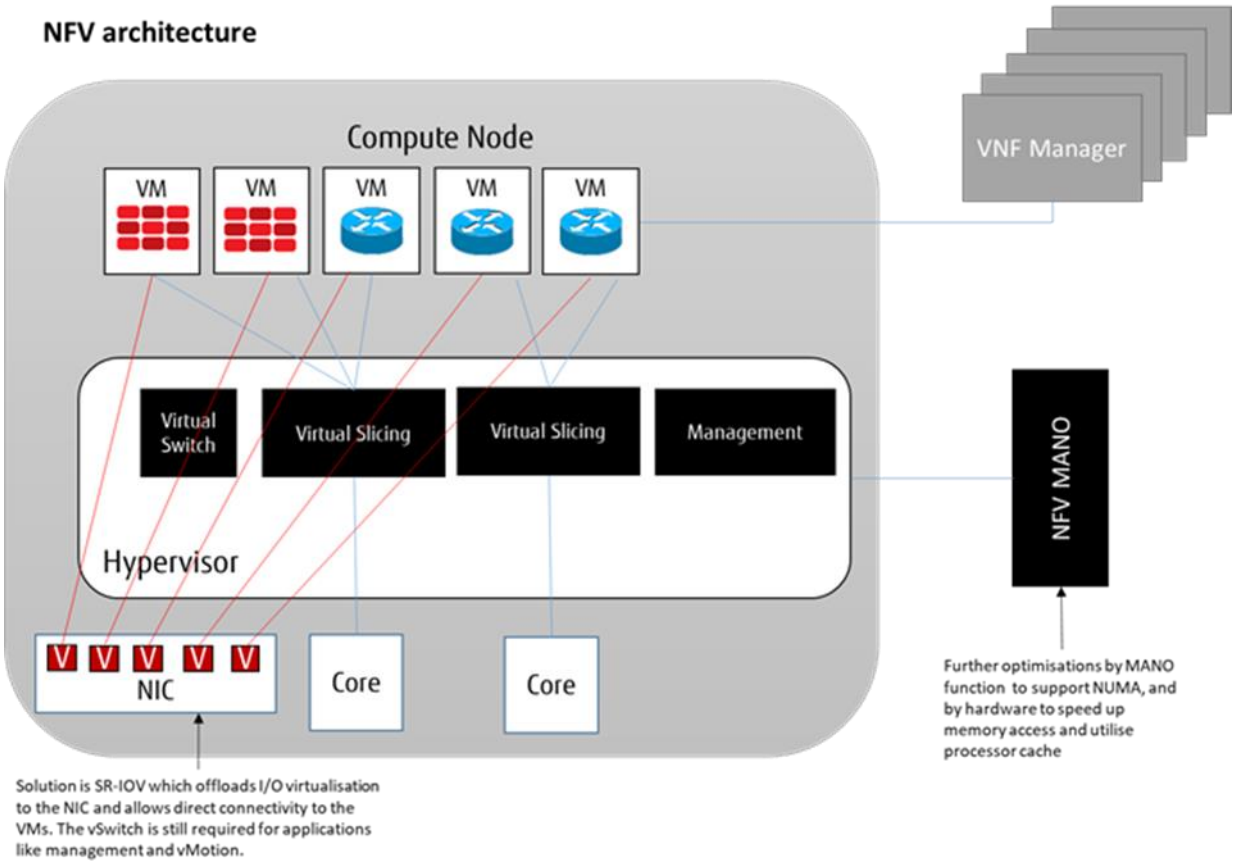

Figure 2

The following limitations apply to the Fujitsu edge devices

Typical Size	Capacity
Connectivity Max Capacity up to 0 to 500mb (typical)	8xCPU 16GB RAM (note SD-WAN will use 4CPU)
Connectivity Max Capacity up to 1GB	16xCPU 32GB RAM (note SD-WAN will use 4CPU)
Connectivity Max Capacity 10GB	32xCPU 64GB RAM (note SD-WAN will use 4CPU)

Table 1 Edge Devices

2. SD-WAN Vendor Edge Devices,

Where the Customer has chosen not to deploy applications in the edge device, use of the SD-WAN vendor manufactured edge device may be more cost effective. Note the on-boarding process and presales activity will vary Fujitsu will advise the Customer of the options available.

Service Operation

Fujitsu's SD-WAN service comprises of Orchestrator tooling, Edge device or software. The platforms are supported by Fujitsu. Unless otherwise detailed in the Statement of Work, the Customer is responsible for the day-to-day management of the service, policies and configuration of the edge devices deployed, with access to Fujitsu professional services to support at the SFIA rate card. Fujitsu will be responsible for the availability of the core infrastructure and availability of the Orchestration Platform. Faults caused by the Customer are excluded from Fujitsu Service Level performance.

Using the (optional) Premium Service the SD-WAN Service will be fully managed end –to-end by Fujitsu (aligned to the Customer Information Assurance policies), or a combination of Customer and Fujitsu performing agreed roles and responsibilities. Where required, Fujitsu will agree the scope of requirements and design a service management solution using the appropriate SFIA rate card charges and management charges.

Fujitsu network support and administration teams will monitor and manage the system and perform support and maintenance tasks. We can offer these in two service levels:

1. Standard Service

Where the Customer choose to manage the SD-WAN Service. Fujitsu analysts will monitor the capacity and health of the infrastructure on a 24/7 basis. Customers will receive support from the Fujitsu support team between 9am-5:00pm, Monday-Friday excluding Bank Holidays, to provide assistance when health, capacity or other issues are identified. Where onsite edge devices prove to be faulty Fujitsu will dispatch next day replacement to site. Faults may be reported on a 24/7 basis. Various Service uplifts are available for access to a Fujitsu service desk and assistance to 24/7.

2. Premium - Full Administration (optional)

This is for Customers who require Fujitsu to manage all aspects of the SD-WAN Service. Fujitsu will provide a fully managed service covering Moves, Adds, Changes and Deletions (MACD) and SD-WAN policies. The number of MACD requests processed will be limited by a ceiling, which shall be agreed with the Customer and recorded in the Customer order form (from the Statement of Work). Volumes in excess of the ceiling will be outside the agreed SLA and subject to extra charges. This is a 24/7 service offering.

Support Services

The Standard Support includes the following features:

- Helpdesk service (in contracted hours)
- Option for Helpdesk 24x7
- Customer portal
- General encrypted overlay product assistance (not configuration changes)
- Provision of documentation, Software upgrades and patches
- Replacement of failed edge devices (Appliance next business day, NFV upon request).

The Premium Support includes the following features:

- Helpdesk service (Available 24x7)
- Online Customer portal

- SD-WAN management including creation and management of routing policies
- Provision of documentation, Software upgrades and patches
- Creation of Encryption Policies and scripts
- Replacement of failed edge devices (Appliance next business day, NFV upon request).

Fujitsu uses its team of SD-WAN specialists to provide Customers with the highest levels of support. The team has many years of experience and holds the full range of Fujitsu and vendor certifications. This service includes hardware (if applicable) and software maintenance along with access to the vendor 4th line support (account teams and development teams), in addition to the Fujitsu provided services.

SD-WAN Professional Services

As a specialist in SD-WAN services, Fujitsu has worked with a number of Public Sector Customers to whom it delivers overlay solutions. With expert knowledge in NFV and SDN, Fujitsu has successfully implemented solutions to fulfil the requirements for NFV, SDN, SD-WAN and emerging SD-LAN services. These solutions are secure, highly available and compliant against Customer Information Assurance policies. Fujitsu's consultants are certified specialists and hold appropriate UK security clearances, including access to DV cleared personnel.

Fujitsu can provide all aspects of SD-WAN overlay Professional Services delivery including:

- Requirement gathering assistance
- Project Management
- SD-WAN overlay Architecture Design
- SD-WAN overlay Deployment
- SD-WAN overlay Migration Assistance
- All services required are chargeable in accordance with the SFIA catalogue.

2. SD-WAN for the Law Enforcement Community, Service Pack 3

Service Pack 3 has been developed as a G Cloud SD-WAN service for the Law Enforcement Community (LEC) operating under the PDS / NPIRMT / PASF Assurance schemes. The description below provides a detailed summary of the functionality of the service and optional services available as standard. Service Pack 3 comprises of a Fujitsu's SD-WAN Service using Cisco (Viptela) technology and complementing applications. The SaaS solution comprises of service management options calculated using the SFIA rate card. We can provide the SD-WAN licences required to support the SD-WAN functionality on a rental basis, based on the term committed. Where a Customer has sufficient Cisco SD-WAN transferable licenses, they can be deployed on the Service Pack 3 platform subject to validation of license terms (and a pro rata adjustment to charges)

2.1 Purpose of Document

2.1.1 This section (2) describes Service Pack 3 SD-WAN service comprising of:

(a) SD-WAN network solution overview

(b) Infrastructure provided to deliver the service, including management toolsets and Customer edge devices.

(c) SD-WAN overlay and management connectivity.

(d) Underlay network dependencies

(e) Network security

2.2 Risks, Assumptions, Issues and Dependencies

2.2.1 A Law Enforcement Community SD-WAN service RAID log is maintained that consolidates internal and external factors that affect the security of the solution described in this design, this is a separate document, maintained by Fujitsu on behalf of the user community.

2.3 Constraints (Standards, Policies, Guidelines)

2.3.1 The solution described will be delivered in accordance with NCSC guidelines.

2.3.2 Aligns to PDS / NPIRMT / PASF assurance

2.3.3 The known guidance documents are listed below:

(a) Network Security

(b) Using IPSec to Protect Data

(c) Using TLS to Protect Data

(d) Cloud Security Guidance

(e) Customer provided Information Assurance standards

2.3.4 Service Pack 3 for LEC is delivered in accordance with the Customer Authority Security Aspects Letter (SAL).

2.3.5 Service Pack 3 comprising of an SD WAN service enables Customers to manage or out-task and outsource the management of a Cisco SD-WAN operating over Customer provided network connectivity, and is called overlay network. For the LEC network connectivity comprises of the PSNfP / PSN Assured (sun setting), a private MPLS network, Internet/ISP connectivity or 5g (ISP) and is called the underlay network).

2.3.6 The SD WAN baseline service includes SD-WAN network design, delivery, project management, SD-WAN network maintenance and SD-WAN network security services using the SFIA rate card.

2.3.7 The SD-WAN service is comprised of the following main components:

(a) **Core Platform** – Orchestration, management and control using Cisco SD-WAN vManage, vBond and vSmart appliances. The core platform is operated with high availability in centralised compute infrastructure.

(b) **Edge devices** – The SD-WAN Edge devices are deployed on Customer sites, private clouds and public clouds.

(c) **SD-WAN Overlay** – The SD-WAN overlay provides encrypted data path connectivity between SD-WAN Edge routers. The SD-WAN Overlay is transported over the Customer provided underlay network infrastructure.

(d) **Analytics portal** – Non-Cloud based analytics provide an overview of SD-WAN network performance over time. This functionality is supported using Live NX hosted in Fujitsu DNS infrastructure

2.3.8 The SD-WAN baseline service:

(a) Manages SD-WAN components of the solution and their connectivity to a number of transport networks (the underlay networks).

(b) Maintains secure management and control plane connectivity between the SD-WAN edge devices and the core platform.

(c) Maintains the set of Cisco standard encryption tunnels built between the SD-WAN edge routers as required by the agreed SD-WAN overlay network topology.

(d) Maintains the set of encrypted Security Associations between the edge devices as required by the overlay network topology.

2.3.9 The SD-WAN baseline service monitors but does not maintain or operate the various underlay networks that the SD-WAN overlay is using for connectivity.

2.3.10 Where a failure in the SD-WAN overlay network is caused by a component of the underlay network Fujitsu will inform the Customer (ITIL aligned procedures and ITSM trouble ticketing) so that the Customer can resolve the issue with their network supplier(s).

2.3.11 The creation of routing policies and application performance requirements shall be the responsibility of the Customer or Authority to test design and provide

2.3.12 The standard SD WAN service demarcation is shown below.

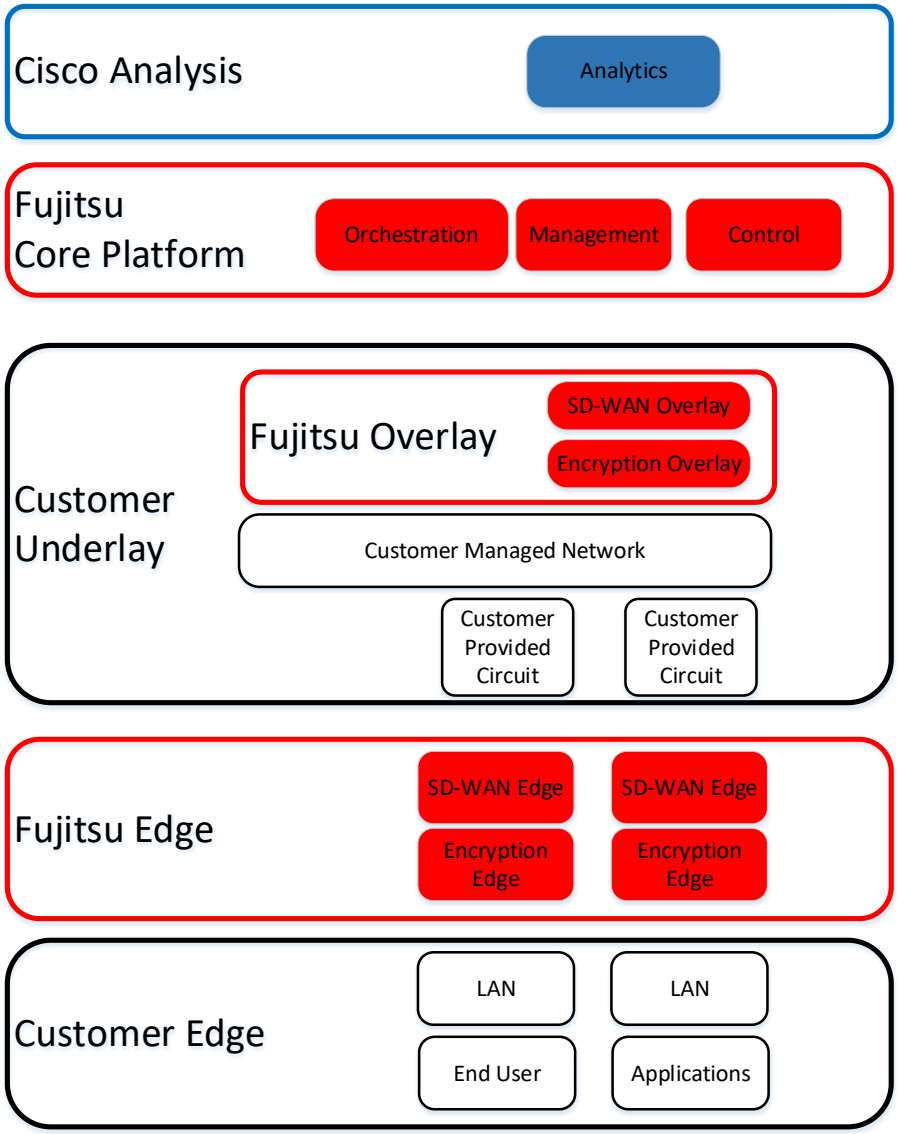

Figure 1 - Standard SD-WAN Service Boundaries

2.3.13 Where the SD-WAN management and control plane is running in Fujitsu infrastructure then Fujitsu shall also be responsible for incident and problem management for that infrastructure to ensure the control plane is available within the committed Service Level Agreement.

2.3.14 The SD-WAN baseline service provides the following service management capabilities, based on the SFIA rate card:

(a) Overlay Network Monitoring (service element 1)

(b) Incident and Problem analysis (service element 2)

(c) Break-fix management (service element 3)

(d) Service Reporting (service element 4)

(e) Patch Management (service element 5)

(f) Change Assessment (Service Element 6)

2.3.15 The G-Cloud SD WAN service demarcation is shown below.

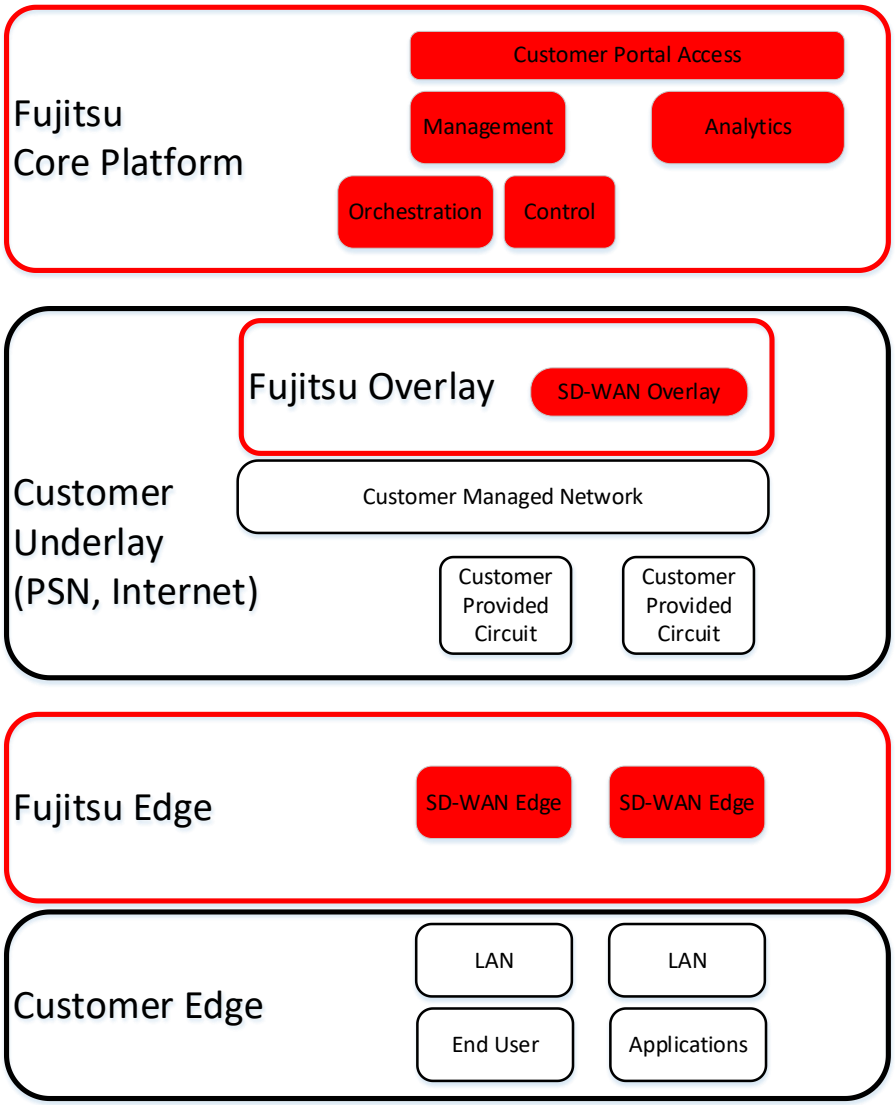

Figure 2 - SD-WAN Service Boundaries

2.3.16 The SD-WAN solution combines Cisco SD WAN capabilities, as published by Fujitsu (functionality guide) with Network Function Virtualization (NFV) in a single Universal Customer Premises Equipment (uCPE), as illustrated in the following figure.

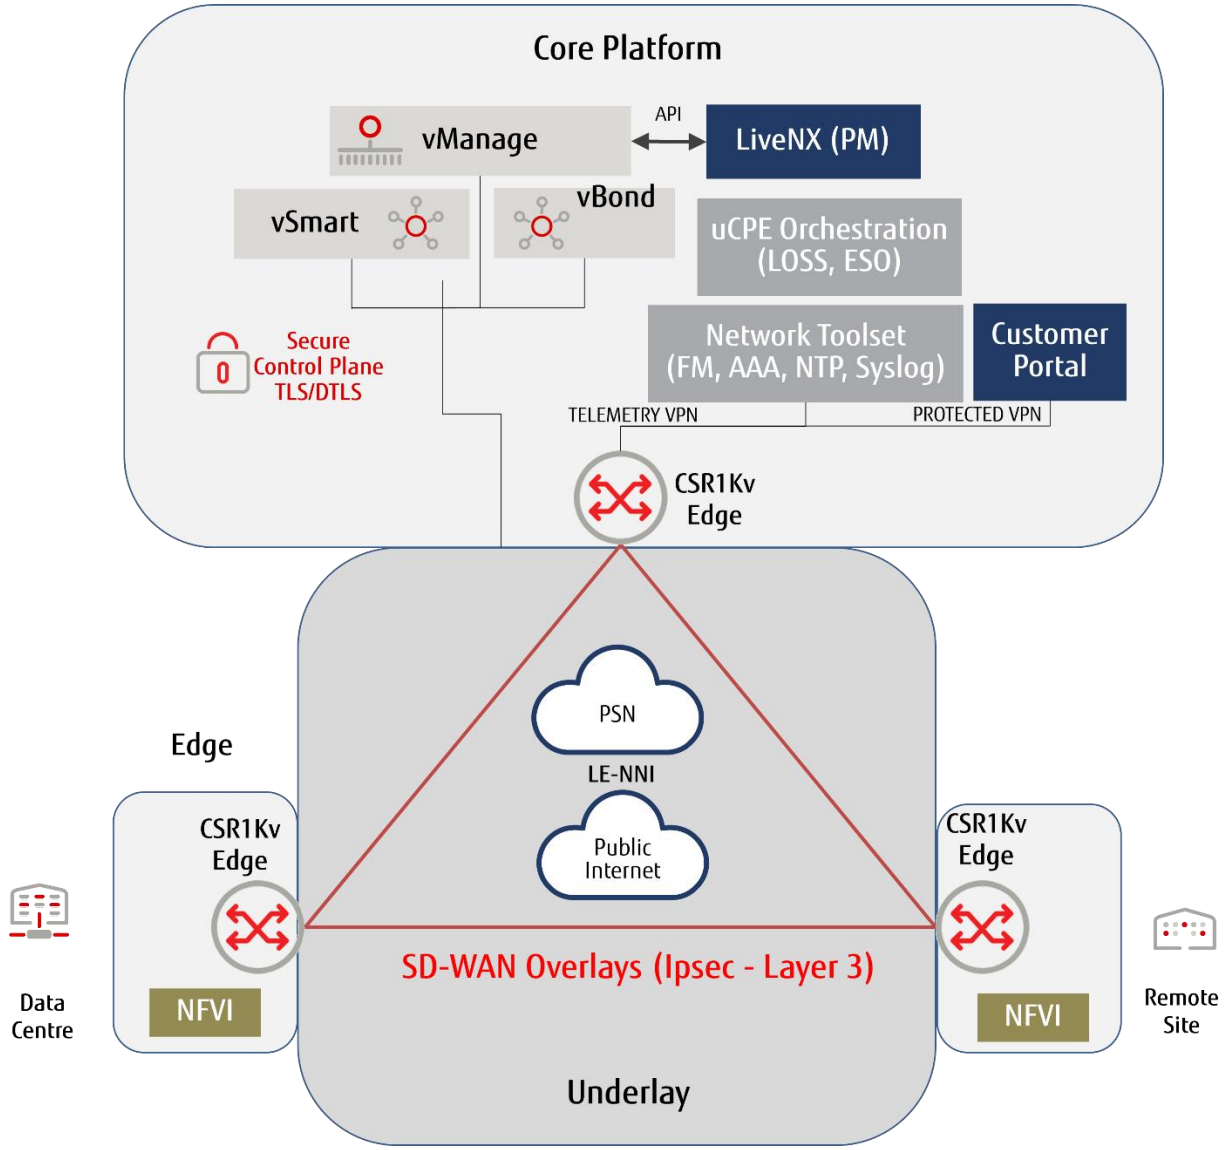

Figure 3 - SD-WAN Architecture

2.3.17 The Core Platform provides the SD WAN service orchestration, management and control functions and includes the following main solution components:

(a) Edge uCPE orchestration

(i) Fujitsu Lightweight Operation Support System (LOSS) enables secure near zero-touch on-boarding of uCPE at Customer sites. Adva Ensemble Service Orchestration (ESO) completes the on-boarding process with secure VNF service chain orchestration.

(b) SD-WAN overlay management and control

(i) Cisco vManage is a centralised dashboard that facilitates automatic configuration, management and monitoring of the SD WAN overlay network and SD WAN Edge routers. Users login to vManage to centrally manage all aspects of the network life-cycle from initial deployment, ongoing monitoring and troubleshooting to change control and software upgrades.

(ii) Cisco vSmart Controllers establish Secure Socket Layer (SSL) connections to all other components in the SD WAN network. They also run an Overlay Management Protocol (OMP) to exchange routing, security, and policy information. The centralised policy engine in vSmart Controllers provides policy constructs to manipulate routing information, access control, segmentation, extranets and service chaining.

(iii) Cisco vBond Orchestrator facilitates the initial SD WAN router bring-up by performing authentication and authorisation of all elements into the network. Cisco vBond Orchestrator also provides the information on how each of the components connects to other components.

(c) Network Performance

(i) LiveNX performance management tool provides network and application analysis and reporting. The LiveNX appliance supports application monitoring and analysis.

(d) Customer portals

(i) Customer Portals are provided for the Cisco vManage and LiveNX management appliances. Remote access is provided to authorised users. Access to the portals is authorised by Role based Access Control (RBAC).

(ii) The Customer Portal provides web server access via a proxy service on the SD WAN overlay to Customers.

(iii) The Customer Portal does not provide a Remote Access Service for access via the internet.

(e) Infrastructure and enterprise management

(i) Network Toolsets complete the management solution, with Zabbix event management, LiveNX performance management, real-time clock provision with Network Timing Protocol (NTP), Elasticsearch syslog collectors, private Domain Name System (DNS) servers, Public Key Infrastructure (PKI) services and Role Based Access Control (RBAC) authentication, Veeam backup and recovery services (BAR), file import/export facilities (IMPEX), download servers, firewalls, intrusion detection/prevention services (IDS/IPS) and anti-virus checking.

(ii) Syslog messages generated by SD-WAN edge devices (C8000v, EC), managers (vManage) and controllers (vBond, vSmart) shall be forwarded to a nominated central customer syslog server from the Elasticsearch syslog collector via the SD-WAN overlay.

2.3.18 The Edge Devices provide Network Function Virtualization Infrastructure (NFVI) to host SD-WAN and encryption virtual edge appliances:

(a) The edge devices are universal Customer Premises Equipment (uCPE) providing compute, memory and network resources using commercial, off-the-shelf (COTS) servers.

(b) The uCPE is available with both single and dual power supply options to allow for value for money at standard sites and higher availability at critical sites.

(c) uCPE devices can be deployed in both non-resilient and high availability configurations, with warm standby units provided at non-resilient sites to minimise service interruption in the event of a device failure.

(d) uCPE devices are pre-staged with the (NFVI) and a basic factory default SD–WAN edge router with WAN/LAN port connectivity in Fujitsu's secure pre-staging facility.

(e) Once the uCPE has been installed on-site, using Fujitsu engineering or smart hands, it is on-boarded to the SD-WAN manager and controller using Fujitsu's Virtual Edge on-boarding process.

(f) After validation of the site installation, the SD-WAN edge is admitted to the SD-WAN overlay.

(g) Once connected to the LEC toolset, the NFVI can be managed remotely from the ESO/LOSS toolset to modify the uCPE and SD-WAN service chain.

2.3.19 The SD-WAN overlay provides intent-based routing of the Customers applications over segmented virtual private networks (VPNs) using IPsec encrypted tunnels:

(a) The SD-WAN overlay uses the underlay and internet transport networks.

(b) The SD-WAN overlay provides network connectivity between SD-WAN edge routers.

(c) The overlay protects data in transit using IPsec tunnels with PKI certificate based authentication, AES-256 encryption and centralised key management.

(d) The SD WAN Edge is deployed as a C8000v virtual router on the uCPE hosting platform.

(e) C8000v Edge routers are located at the Customer site or in the cloud and provide data plane connectivity between sites via the SD-WAN overlay.

(f) The edge routers provide packet forwarding decisions, policies and processes at each site in conjunction with the centralised vSmart controllers.

(g) SD-WAN provides centralised dynamic routing using the vSmart controller and the Overlay Management Protocol (OMP).

(h) The SD-WAN edge combines proactive overlay performance monitoring with application based routing policies to enable optimisation of network path selection.

(i) Policies can be independently defined for applications, application groups, network segments (VRFs) and sites to meet different organisational and business needs.

(j) Cflowd policies can be configured to gather and export detailed application flow information for analysis in the vManage and LiveNX management appliances.

(k) The SD-WAN edge router includes a stateful enterprise firewall with application awareness.

2.4 Core Platform

2.4.1 Orchestration, Management and Control Architecture

(a) The Core Platform provides orchestration, management and control of the integrated uCPE, SD WAN and analysis functions in a high availability (HA) architecture, as illustrated in the following figure.

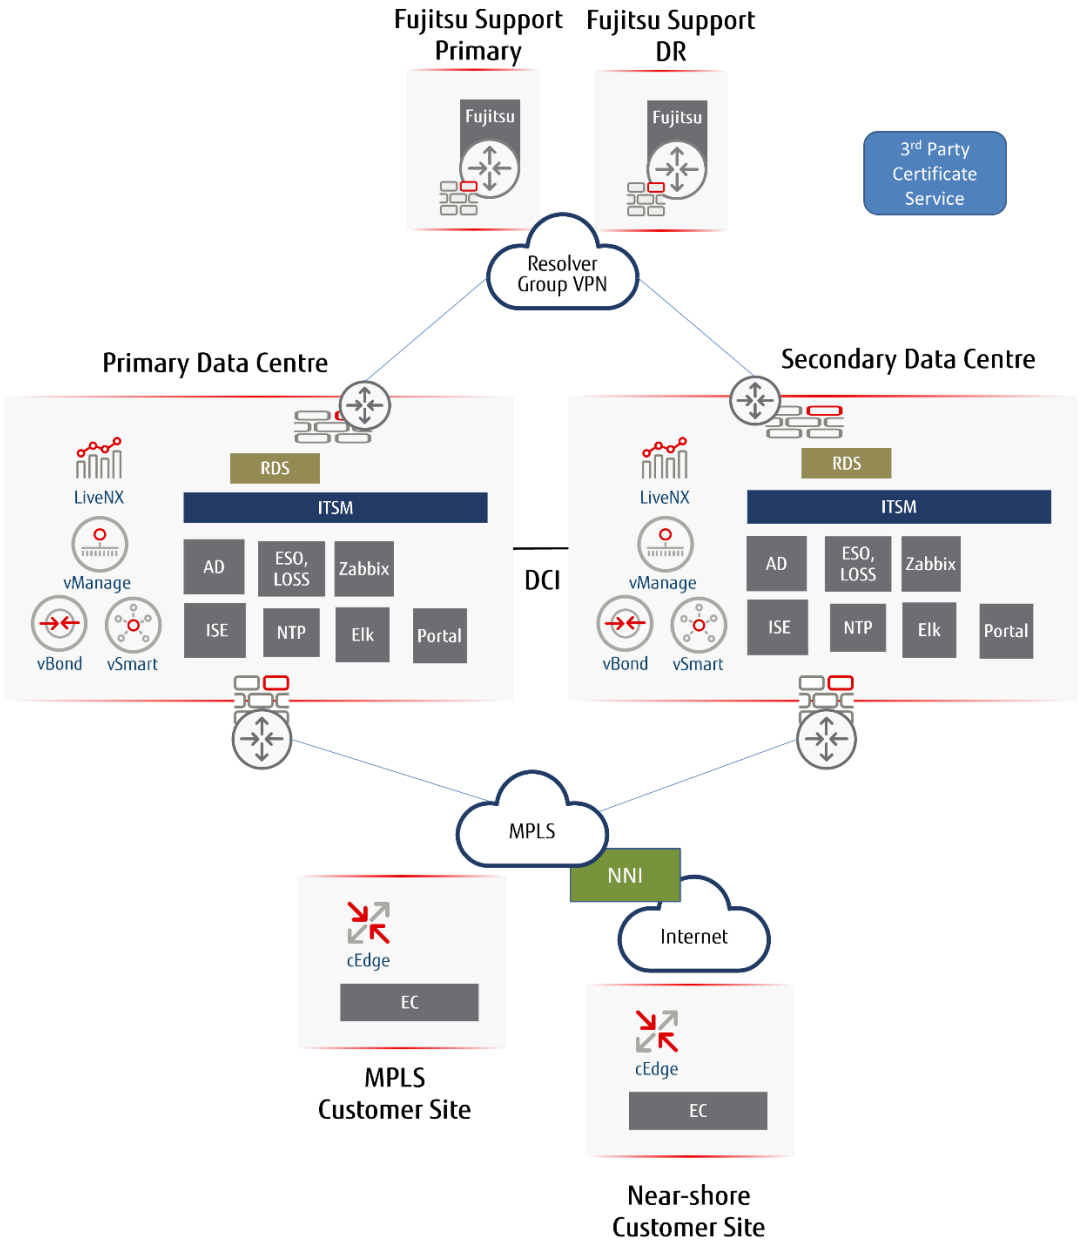

Figure 4 - High Availability Toolset.

(b) The primary and secondary toolsets will be hosted on the Service platform operated at an OFFICIAL-SENSITIVE handling caveat.

(c) The toolset will enforce secure protection of the Customer data through use of secure management protocols, network and transport layer encryption, encryption of sensitive data at rest, security event logging and role-based access control.

(d) The SD-WAN toolset will be deployed in a High Availability architecture, with each management, orchestration and controller appliance protected from failure by a backup appliance in the geo-separated secondary datacentre.

(e) Individual appliances will be operated in active/active or active/standby mode, depending on the appliance and the service provided.

2.4.2 Toolsets

(a) The LEC SD-WAN toolset components of a single datacentre are illustrated in the following figure.

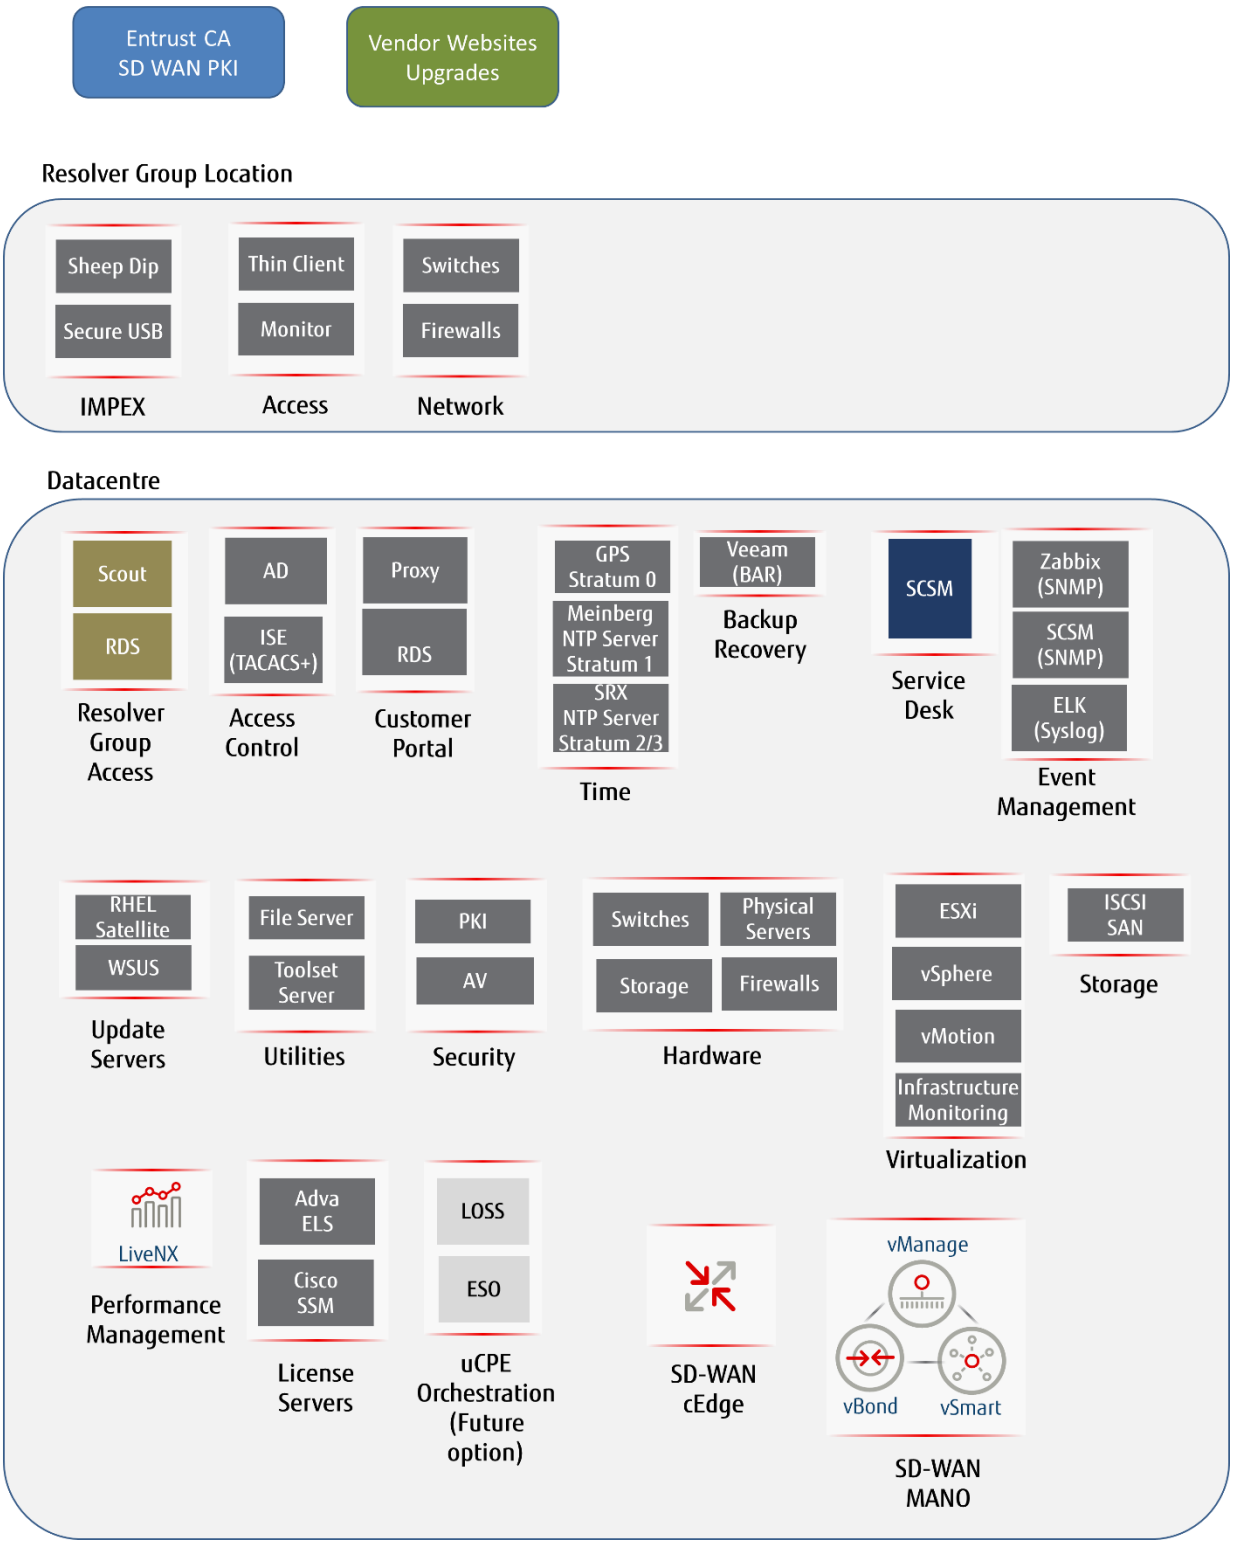

Figure 5 - SD-WAN Toolset Components

(b) The toolset is split into shared tools, providing management support for multiple Customers operating at the same security classification, and dedicated tools supporting a single Customer, as illustrated in the following figure.

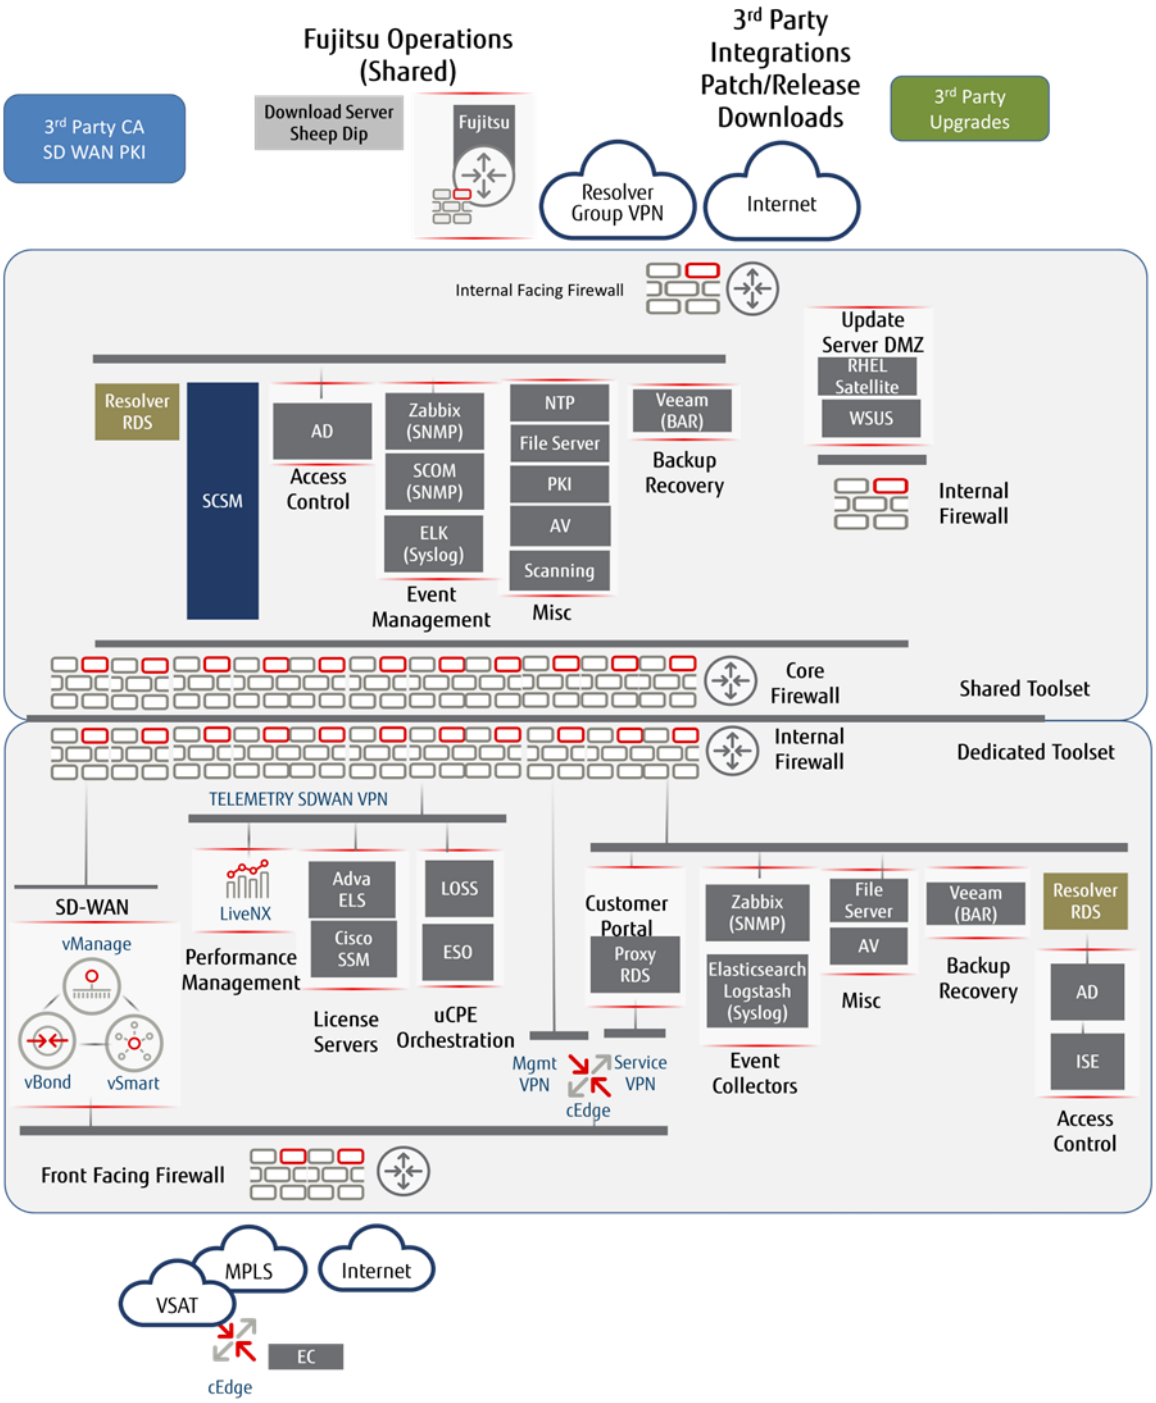

Figure 6 Dedicated and Shared Toolsets

(a) The following toolset appliances and applications are shared with other Fujitsu network Customers:

(i) Windows System Centre Services Manager (SCSM)

(ii) Windows file server

(iii) Windows Server Update Services (WSUS)

(iv) Redhat Linux (RHEL) satellite server

(v) Mcafee anti-virus

(vi) Windows Access Directory (AD)

(vii) Meinberg NTP time servers with GPS time source.

(viii) SCOUT terminal server and shared Remote Desktop Servers (RDS).

(ix) Internal PKI CA

(x) Keepass password manager

(xi) Sheep dip with Symantec anti-virus

(xii) Zabbix

(xiii) Elasticsearch, Kibana and Logstash (ELK) stack

(xiv) Veeam

(xv) VMware vSphere (ESXi and vCentre)

(b) The shared toolset supports the following functionality:

(i) ***Event Management***: Separate Zabbix proxies are provided to collect event notifications in each dedicated Customer zone. The Zabbix servers in the shared zone provide management access to dedicated Customer domains, event correlation capabilities and automated SCSM ticket generation.

(ii) ***Security event management***: Separate Elasticsearch collectors are provided in each dedicated Customer zone. In the shared zone, an Elasticsearch, Kibana and Logstash (ELK) stack provides management access to dedicated Customer domains and event correlation capabilities. ELK is integrated with Zabbix in the shared zone to provide automated SCSM ticket generation.

(iii) ***Internal Public Key Infrastructure:*** The toolset provides PKI for web services and internal encryption services.

(iv) ***Access Control:*** Microsoft Access Directory (AD) provides domain control in the shared zone with role-based access control (RBAC). The shared AD domain controller manages access to the dedicated zones.

(v) ***Fujitsu Toolset access:*** The shared zone provides SCOUT terminal servers to enable remote access from resolver group thin clients. The shared Remote Desktop Server (RDS) provides access to the shared tools and access to the dedicated Customer zones via a separate RDS in each zone.

(vi) ***Password management***: The shared Keepass generates and stores complex passwords for local accounts.

(vii) ***Real Time Clock***: All network devices and toolset appliances are locked to a common Universal Time Co-ordinated (UTC) source, provided by the toolset Network Timing Protocol (NTP) server. The NTP is locked to a Global Navigation Satellite System (GNSS) stratum 0 clock and provides a stratum 1 timing source.

(viii) ***Patching:*** WSUS and RHEL satellite servers provide automated solution for identifying and downloading patches for Windows and Linux servers.

(ix) ***Import/Export:*** Download server and sheep-dip antivirus servers are provided to facilitate secure import and export of files, including software images and reports. Files are downloaded, checked for viruses and malware and copied to a secure USB for transfer to the thin client and download to the shared toolset file server.

(x) ***Backup and recovery:*** Veeam is used to provide automated file and VM server backup for replication and disaster recovery.

(xi) ***Virtualization Infrastructure:*** VMware vSphere provides the datacentre virtualization platform.

(c) The following toolset appliances are dedicated to Service Pack 3 SD-WAN users:

(i) SD WAN: vManage, vBond, vSmart

(ii) Performance Management: LiveNX

(iii) uCPE orchestration: ESO, LOSS

(iv) Syslog collector: Elastic Stack

(v) Event management collector: Zabbix

(vi) Keepass password management

(d) Note: The SD WAN vManage, vBond, vSmart and LiveNx applications are dedicated for LEC (per user group, such as [redacted]) in order to support sensitive configuration and application data relating to the network and user applications.

(e) The dedicated toolset supports the following functional groups:

(i) ***uCPE Orchestration***: The Fujitsu Lightweight Operation Support System (LOSS) and Adva Ensemble System Orchestrator (ESO) provide orchestration of the uCPE, the Kernel-based Virtual Machine (KVM) hypervisor (Edge Connect), Virtual Network Functions (VNFs) and service chaining. LOSS and ESO are only used during service chain orchestration.

(ii) ***SD WAN***: The SD-WAN toolset includes the Cisco SD-WAN orchestrator (vBond), controller (vSmart) and manager (vManage). vManage provides single pane of glass management of Fujitsu's SD-WAN solution and include REST API for integration with other toolsets.

(iii) ***Event Collector***: Edge device events and notifications are forwarded to the Zabbix event management collector using SNMP traps.

(iv) ***Security Event Collector:*** Syslog messages are reported to the dedicated Elasticsearch collector. The Elasticsearch syslog collector relays messages to the Customer's Security Incident and Event Management (SIEM) platform via the SD-WAN overlay.

(v) ***Performance Management***: The LiveAction performance management solution (LiveNX) is fully integrated with vManage and provides network Service Level Agreement (SLA) reporting and application analysis.

(vi) ***Public Key Infrastructure:*** The SD-WAN solution uses a 3rd party certificate signing service (Entrust Platinum) as the root certificate authority (root CA), with the SD-WAN vManage acting as the PKI root CA for the edge devices.

(vii) ***Password management***: The dedicated Keepass generates and stores complex passwords for local accounts in the dedicated toolset and remote network devices.

(viii) ***Customer Portal Access:*** A dedicated windows server proxy provides secure user access to the LEC dedicated SD-WAN and LiveNX toolsets.

2.4.3 Management Connectivity

(a) The following table summarises the LEC SD-WAN toolset management and control connectivity.

Remote Function	Toolset Device	Secure Communication Channel	Datacentre Terminating Network Device	Transport Network
Fujitsu Resolver Groups	Thin client and Remote Desktop Service	[redacted]	[redacted]	[redacted]
Customer Portal	Proxy and Remote Desktop Server	[redacted]	SDWAN cEdge (C8000v) (Dedicated)	MPLS / PSN / NNI SDWAN overlay – Service VPN
Customer SIEM	Elasticsearch syslog collector	[redacted]	SDWAN cEdge (C8000v) (Dedicated)	[redacted]
vManage, vBond, vSmart	vBond	[redacted]	vBond	MPLS / PSN/Internet underlay
vManage, vSmart	vManage, vSmart	[redacted]	vManage, vSmart	MPLS / PSN/Internet underlay
cEdge (C8000v)	vManage, vSmart	[redacted]	vManage, vSmart (Dedicated)	MPLS / PSN/Internet underlay
cEdge (C8000v)	vBond	[redacted]	vBond (Dedicated)	
	NTP server	[redacted]	cEdge (C8000v) (Dedicated)	MPLS / PSN/Internet SD-WAN overlay – Telemetry VPN
	Syslog		Syslog collector (Dedicated)	
	LiveNX (Flexible Netflow)		cEdge (C8000v) (Dedicated)	

Remote Function	Toolset Device	Secure Communication Channel	Datacentre Terminating Network Device	Transport Network
	ISE (TACACS+)		cEdge (C8000v) (Frontend)	
	Zabbix (SNMP traps)		cEdge (C8000v) (Frontend)	
	LOSS/ESO		cEdge (C8000v) (Dedicated)	

Table 1 SD-WAN Management and Control Connectivity

(b) NTP communications will be authenticated where possible.

2.5 Edge Devices

2.5.1 Universal CPE

(a) The optional SD-WAN Virtual Edge (shown in the figure below) provided by Fujitsu is a universal Customer Premises Equipment (uCPE) that supports virtual SD WAN Edge routers and third-party Virtual Network Functions (VNFs) at aggregate WAN capacities of up to 4.5 Gbps. The Virtual Edge supports both High Availability and Standard Availability modes of operation.

Figure 7 - uCPE – Dell VEP4600

(b) The Virtual Edge is provided in two variants, the VEP4600 (8 port) and VEP4600 (12 port).

(c) Both the VEP-4600 variants fully support:

(i) 8x WAN/LAN Ethernet ports

(1) The VEP-4600 (8 port) has 4x 10/100/1000baseT RJ45 chassis ports and 4x 10/100/1000baseT RJ45 ports in one expansion slot. The 8-port variant supports 1x WAN port, 3x service VPN LAN ports and 3x service VPN WAN ports for pass-through connectivity, and 1x spare port.

(2) The VEP-4600 (12 port) has 4x 10/100/1000baseT RJ45 chassis ports and 8x 10/100/1000baseT RJ45 ports in two expansion slots. The 12-port variant supports 1x WAN port, 3x service VPN LAN

ports and 3x service VPN WAN ports for pass-through connectivity, and 5xspare ports for additional connectivity.

(ii) Kernel based Virtual Machine (KVM) with Ensemble Connect hypervisor supporting an open, standard interface for VNF service orchestration.

(iii) SD WAN Edge router (VNF)

(d) Both uCPE are provided with the same compute and memory resource:

(i) 8x vCPU

(ii) 16GB RAM

(iii) 240GB SSD storage

(e) The DELL VEP-4600 edge device with dual power supply option.

(f) The DELL VEP 4600 provides

(i) Dual field replaceable power supply modules (1+1 redundancy)

(ii) Field replaceable fan modules (3+1 redundancy)

(iii) 4x RJ45 fixed chassis ports, with a factory fit option for an additional 4x or 8x RJ45 ports

(g) The DELL edge devices and configuration build is deemed to satisfy the requirements of PDS / NPIRMT terminating hardware (end point devices) for encryption and the design deployed is subject to annual ITHC.

(h) Implementation notes:

(i) In order to support patching of the Intel Management Engine (e.g. to fix security vulnerabilities) an enhanced BIOS access level is required to allow remote upgrade.

(ii) For the C8000v supports act/act HA using Virtual Router Redundancy Protocol (VRRP). In this mode traffic is switched between devices via an external WAN/LAN facing switch.

(iii) The end user must provide resilient WAN/LAN interfaces and WAN/LAN facing switches to support HA operation using VRRP.

2.5.2 Edge Device Topology

(a) The LEC universal CPE service chain provides support for the following Virtual Network Functions (VNF) and external port connectivity:

(i) C8000v VNF.

(ii) C8000v VNF external WAN0 Ethernet port connectivity, which is associated with the pre-staged WAN IP configuration.

(iii) Additional C8000v VNF external Ethernet ports, which are assigned to cEdge service VPNs during on-boarding.

(iv) C8000v VNF internal port for connection to the Ensemble Connect host.

(b) Once on-boarded, the SDWAN edge device ports are configured as defined in the following figure.

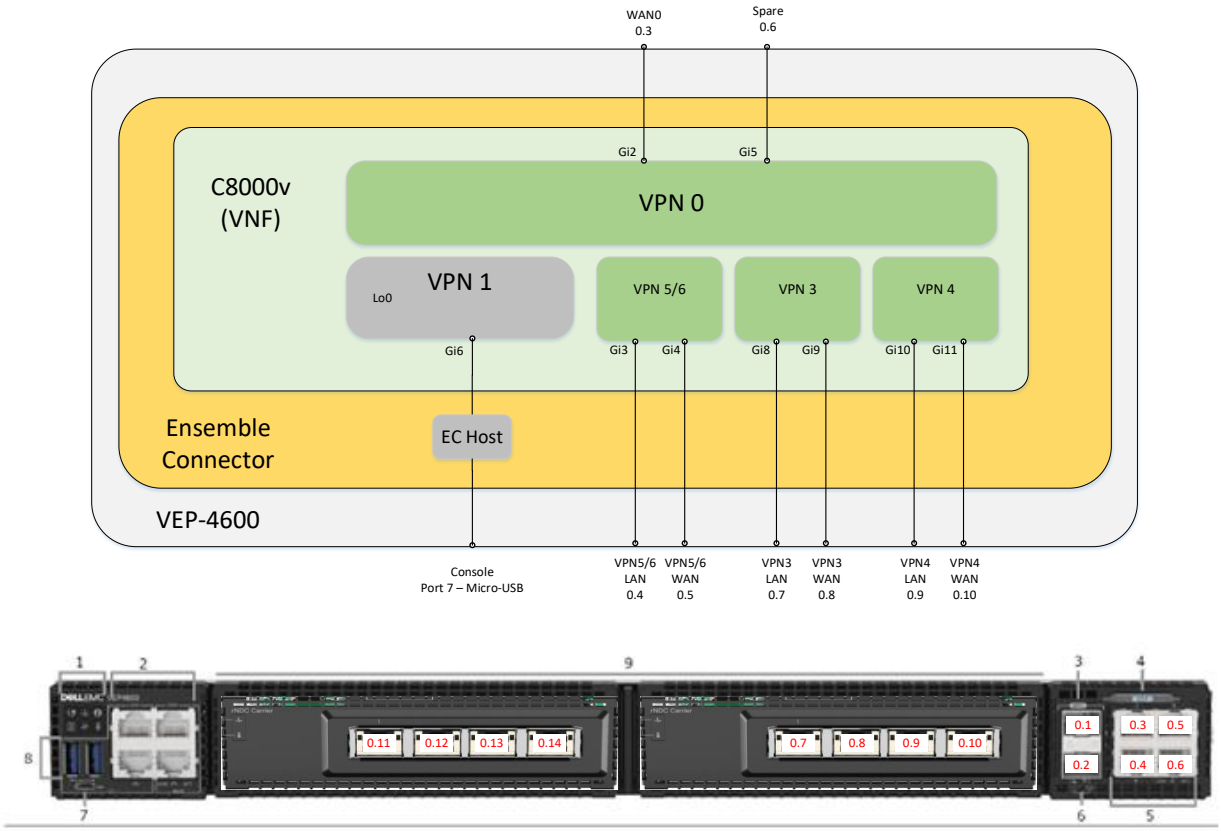

Figure 8 SDWAN Edge – Port Connectivity

(c) The VEP-4600 console port provides access to the Adva EC host.

(d) Local console access to the C8000v cEdge is provided via the Adva EC host and requires separate user authentication.

2.6 SD-WAN Overlay

2.6.1 Overlay Network

(a) The SD-WAN overlay provides IPsec encrypted connectivity between SD-WAN edge devices, as illustrated in the following figure.

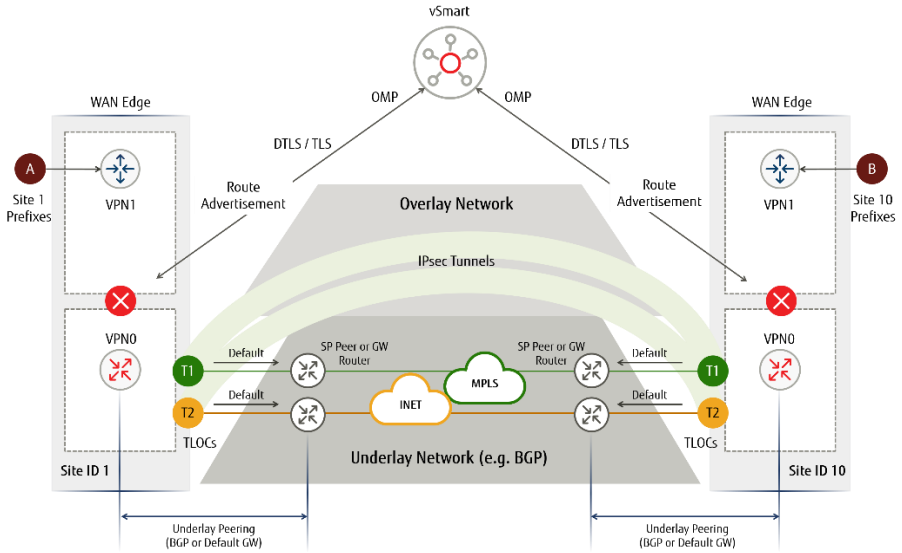

Figure 9 - SD-WAN Overlay Site to Site Connectivity

(a) The SD WAN overlay will provide:

(i) Protection of confidentiality and integrity of Customer traffic [redacted]

(ii) End-to-end segmentation of Customer traffic using service VPNs.

(b) The SD-WAN solution supports different VPN topologies in the SD-WAN overlay to satisfy different Customer connectivity requirements.

(c) VPN templates are used to control the network topologies that define how edge devices are connected via the overlay.

(d) The following overlay topologies are supported:

(i) Full mesh (default)

(ii) Hub and spoke

(iii) Dynamic mesh

(e) The LEC SD-WAN overlay topology will be configured in accordance with the information provided by the Customer in the Cisco Site and Network Configuration & Application Policy Template.

(f) SD-WAN overlay tunnels connected between MPLS / PSN connected sites shall be connected via the MPLS / PSN only.

(g) SD-WAN overlay tunnels connected between internet connected sites shall be connected via the internet only.

(h) The Customer may, using SD-WAN overlay tunnels, connect between MPLS / PSN and Internet connected sites, subject to dialogue with PDS / NPIRMT.

2.6.2 Overlay VPNs

(a) The SD-WAN solution will be operated as a single organisation platform.

(b) All Customers (e.g. police forces and application providers) and Customer networks (e.g. Closed User Groups or service VPNs) will use the SDWAN transport VPN (VPN0) for connectivity of the SDWAN overlay over the PSN/Internet underlay networks. VPN0 is the underlay network.

(c) The SD WAN overlay will provide end-to-end segmentation of Customer traffic, using Virtual Router Functions (VRF) in the service (Customer LAN facing) side of the SD-WAN router and Virtual Private Networks (VPNs) across the SD-WAN overlay.

(d) SD-WAN routers will maintain separate per-VPN routing tables for complete control plane separation.

(e) The SD-WAN overlay VPNs will be configured in accordance with the information provided by the Customer in the Cisco Site and Network Configuration & Application Policy Template.

(f) The SDWAN overlay shall support the following VPNs:

(i) VPN1: SDWAN Telemetry (management).

(ii) Service VPNs as defined by customer

(g) As a default the SDWAN will be operated with the following topologies:

(i) MPLS full mesh: All MPLS connected sites, including NNI edge devices, will be connected in a full mesh for the service VPNs.

(ii) Internet full mesh: All internet connected sites, including the NNI edge devices, will be connected in a full mesh for the service VPNs.

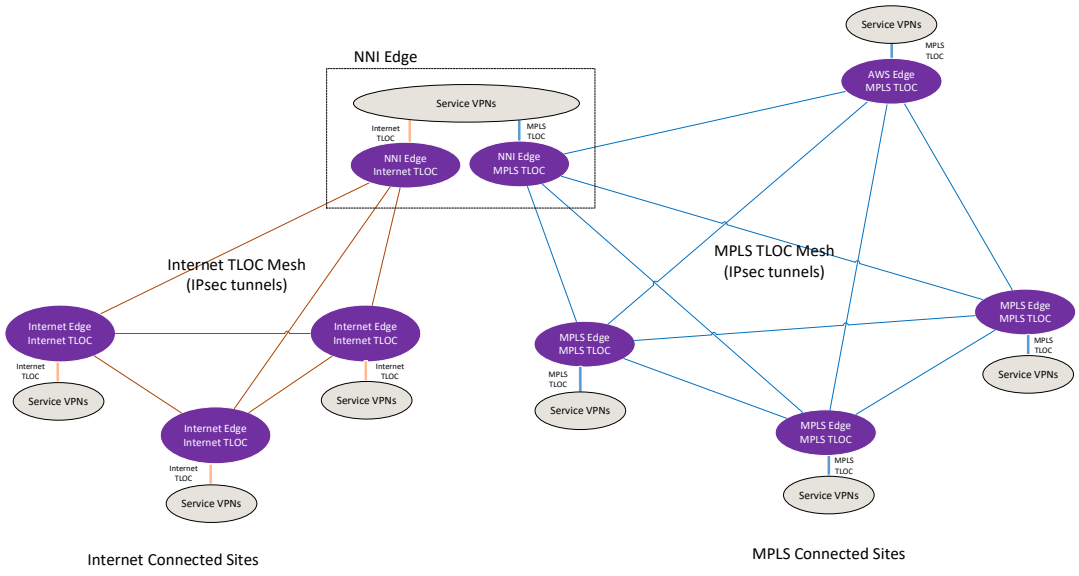

Figure 10 SDWAN Service VPN Topologies

(h) The SDWAN Telemetry VPN (VPN1) is configured in a dual hub and spoke topology, providing resilient datacentre connectivity for each remote edge.

(i) The Syslog VPN (VPN2) is configured with connectivity limited to the SDWAN datacentres and NNI sites.

2.6.3 Dynamic Routing

(a) The SD-WAN overlay supports dynamic routing between SD-WAN edge devices using the vSmart controller and Overlay Management Protocol (OMP).

(b) Traditional dynamic routing protocols, including OSPF and BGP, can also be used in the service VPNs to learn routes in the Customer LAN domain.

(c) vSmart will advertise routes to each SD-WAN edge router in accordance with routing policies.

(d) Routing policies shall be provided to route traffic that is being connected between PSN and internet connected sites via the NNI SDWAN edge.

(e) Control policies will ensure symmetrical routing of traffic across the SDWAN overlay.

2.6.4 Application Quality of Experience

(a) The SD-WAN solution provides a multi-dimensional approach to Application Quality of Experience (AppQoE). An application bias is applied to traditional network functions like QoS and routing. Optimisation features are integrated into the router to supplement the networking functions.

(i) ***Application-Aware Routing***: This dynamically selects paths for specific application flows and groups based on SLA policies and real time network performance. It proactively manages application connectivity, replacing the old paradigm of selecting the most direct network path and hoping for the best.

(ii) ***Quality of Service (QoS):*** QoS provides the basic traffic management functionality that allows application prioritisation. QoS includes traditional classification, policing, scheduling and shaping. AppQoE applies intelligence to the use of these facilities.

(iii) ***Forward Error Correction (FEC)***: FEC and packet duplication features are used to remediate loss over poor quality circuits.

(iv) ***Application optimisation:*** In addition to selecting the optimal path for an application, the SD-WAN solution also includes TCP optimisation. TCP optimisation fine-tunes the processing of TCP data traffic to decrease round-trip latency and improve throughput.

(v) ***Network connectivity***: Multiple high-speed transport network connections will be used to reduce latency and to provide path diversity.

(vi) ***Network*** analytics: LiveNX virtual appliance supports individual network path and application flow monitoring. This will provide a rich data set that enables network diagnostics to proactively manage end user experience.

Application analytics and reporting will identify poorly performing or congested network connectivity. Analysis will reveal how end users are using the network, which enhances focused network planning and capacity management.

(b) The SD WAN service includes capabilities to embrace migration to cloud based services.

(i) ***Cloud onRamp:*** The Fujitsu SD-WAN solution optimises cloud connectivity by dynamically measuring SaaS application performance and selecting the best available path.

(c) Details of the AppQoE features listed above are provided below.

2.6.5 Application Aware Routing

(a) Application Aware Routing enables the ability to dynamically route traffic based on policies and real time connectivity performance (see figure below).

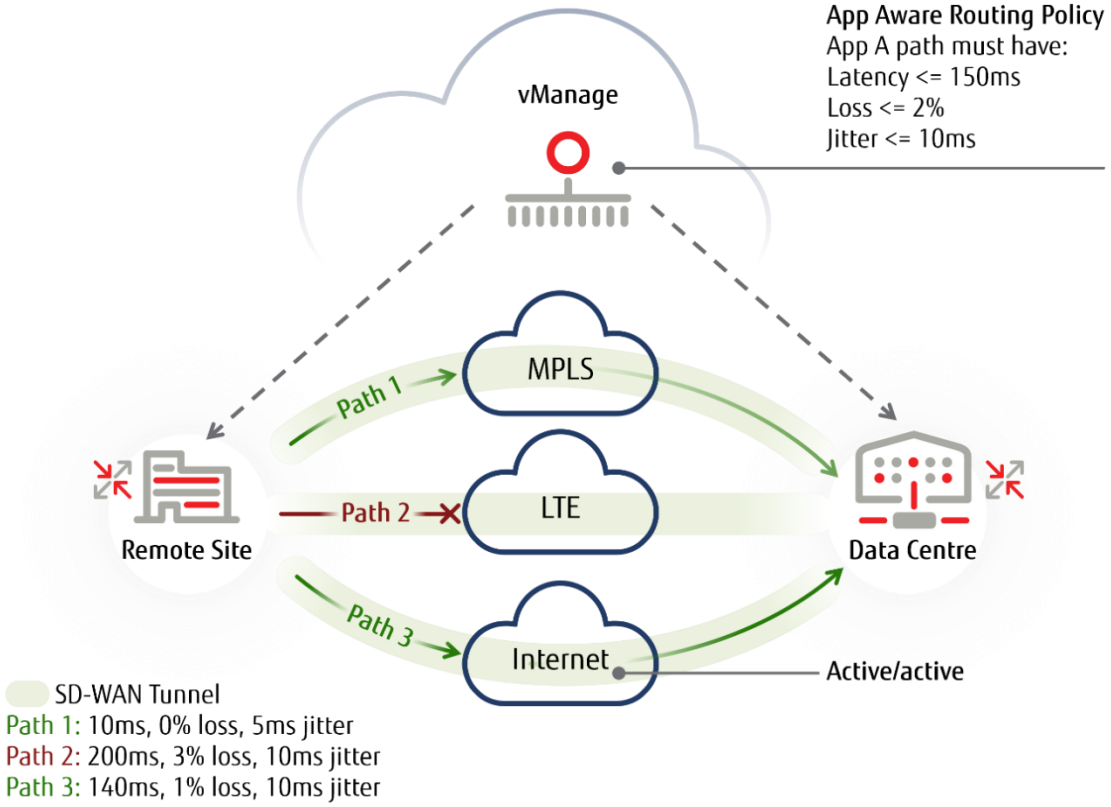

Figure 11 - Application Aware Routing

(b) Connectivity is monitored for latency, packet jitter and packet loss. Latency and jitter metrics enable policies to be defined for real time voice and video applications. Packet loss metrics identify network congestion, allowing applications to be diverted to a different path.

(c) Localised and centralised policies can be defined to match applications and application groups, using one or more of the following criteria:

(i) Source and destination IP address

(ii) DSCP and Packet Loss Priority (PLP)

(iii) Protocol

(iv) Source and destination port

(v) Custom source and destination prefix

(vi) Next Generation Network-Based Application Recognition (NBAR2) application list

(d) The SD-WAN Edge router includes an integrated NBAR2 Deep Packet Inspection (DPI) engine to identify and classify applications including voice and video, email, file sharing, gaming, peer-to-peer (P2P), and cloud-based applications. The NBAR2 DPI engine identifies a wide variety of applications from the network traffic flows using Layer 3 to Layer 7 data.

(e) The SD-WAN Edge Router supports the identification of TLS encrypted applications from Layer 3 (IP and DSCP) and Layer 4 (port and protocol) matching criteria.

(f) Once an application or application group has been defined, a Service Level Agreement (SLA) is configured. The Service Level Agreement specifies the network path characteristics (loss, latency and jitter) that the application can tolerate for optimized performance.

(g) The SD-WAN solution monitors the performance of each network connection continuously, providing real time metrics to the local packet forwarding process and centralised controller. If a link is failing to meet the SLA targets defined for an application, the application will be diverted to the best available path (subject to any service definitions provided by the Authority)

(h) The SD-WAN Edge routers use Bidirectional Forwarding Detection (BFD) to monitor all SD-WAN overlay tunnels. The BFD metrics are used to detect path liveliness and provide loss, latency and jitter quality measurements. The BFD metrics will indicate total loss of connectivity and service degradation.

(i) The SD-WAN solution will use SLA-based policies to choose the optimal path for critical applications and will dynamically switch the path in the event those SLAs are not met.

(j) The LiveNX analytics tool will provide reports based against the SLA policies in near real time. This information can be viewed by the Customers via the portal provided.

(k) These SLA policies define the rules for decisions made by Application-Aware Routing.

2.6.6 SD-WAN Overlay Quality of Service

(a) The SD-WAN routers incorporate the QoS functions listed below, and illustrated in the following Figure:

(i) Input classification

(ii) Input policing

(iii) Output policing

(iv) Output rewriting

(v) Output queuing and scheduling

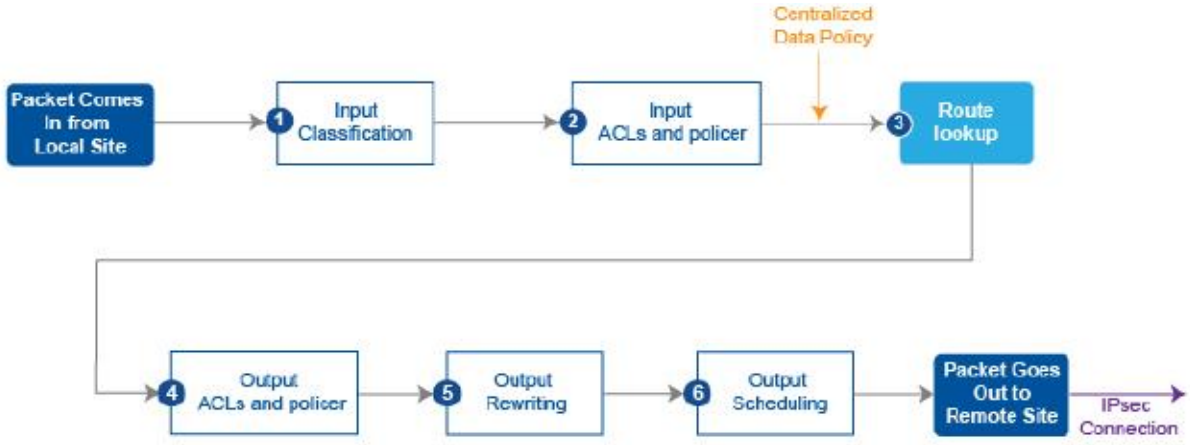

Figure 12 - SD-WAN Edge Router - QoS Components

(b) Input classification maps packets to a defined traffic class, which is used later in the QoS pipeline to decide how a packet is treated.

(c) The SD-WAN solution supports the classification of all packets based on:

(i) Source and destination IP address,

(ii) Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) port numbers,

(iii) Differentiated Services Code Point (DSCP) markings.

(d) This allows the identification of IP conversations and applications by port utilisation, including applications that have been encrypted at Layer 4 using Transport Layer Security (TLS).

(e) Deep Packet Inspection enables additional identification of unencrypted applications by including analysis of Layers 5-7.

(f) The SD-WAN solution can be configured to treat Customer LAN ports as trusted and map the end user DSCP markings to the outer header of the SD-WAN overlay IPsec tunnels, as shown in the Figure below.

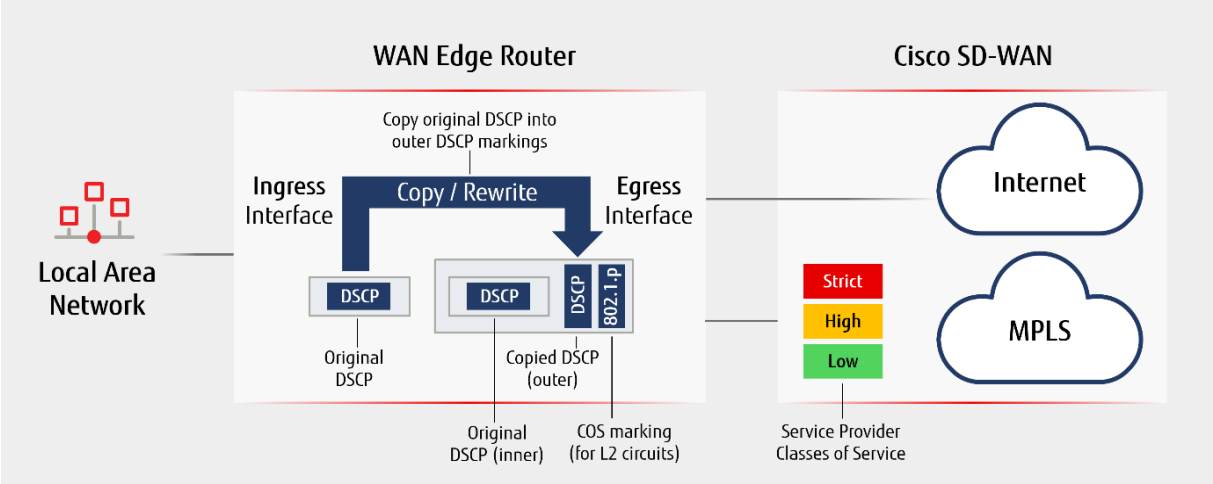

Figure 13 - Overlay Tunnel DSCP Marking

(g) If a port is untrusted, the SD-WAN router will mark the DSCP according to configurable policies. This ensures traffic classes will be honoured in QoS enabled transport networks such as the PSN.

(h) If necessary, the SD-WAN router will re-mark the DSCP to ensure the flow is mapped to the correct traffic class in the transport network.

(i) Queuing and scheduling allows traffic to be transmitted in order of priority at an interface. The SD-WAN router supports eight queues with a low latency queue for real time traffic.

(j) Real-time applications are mapped to high priority traffic classes that ensure priority treatment in the SD-WAN router and QoS enabled transport networks.

(k) The Figure below illustrates the mapping of traffic to strict, high and low priority queues according to application aware policies.

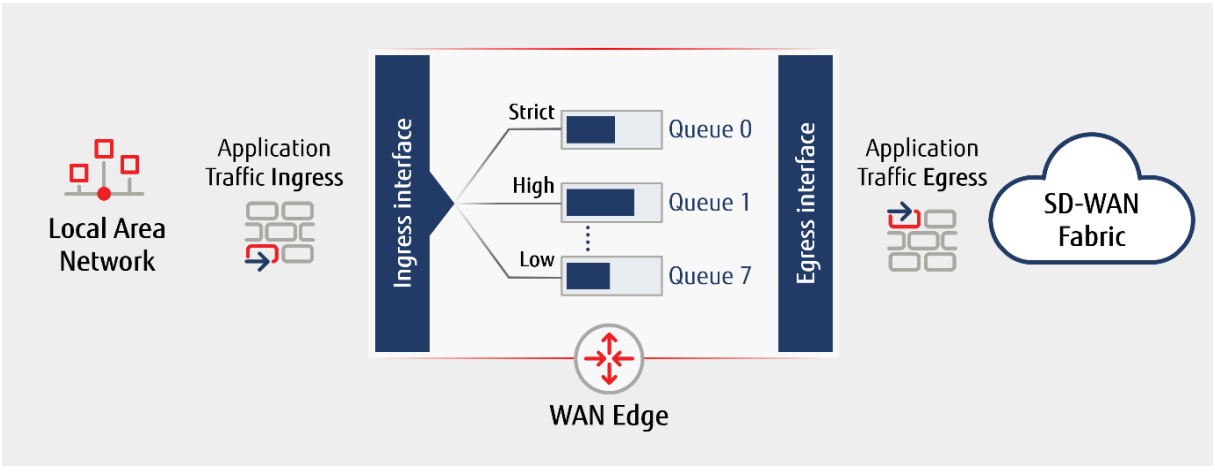

Figure 14 - Application Aware Classification

(l) Input and output policing controls the amount of traffic admitted or transmitted for a specified flow, application or application group. Non-compliant traffic is either re-marked for drop eligibility or dropped, depending on policy.

(m) Shaping controls the maximum rate of traffic transmitted. Traffic will be queued until the shaper determines it can be transmitted in accordance with the defined rate. The SD-WAN router will use aggregate shaping at an interface to maintain compliance with the PSN/Internet WAN interface rate.

(n) The SD-WAN solution supports the use of per-tunnel QoS at the hub sites of a hub and spoke topology to ensure smaller spoke sites are note starved of connectivity by larger sites sharing the same hub.

(o) Per-tunnel QoS will shape the maximum rate of each SD-WAN tunnel, effectively partitioning the hub bandwidth between the spoke sites.

2.6.7 SD-WAN Underlay QoS

(a) Underlay QoS is provided by the PSN or a private MPLS network.

(b) There is no QoS in the internet underlay.

(c) The DSCP marking of GRE and IPsec tunnels that are transported over the WAN transport network will be derived from the end user DSCP marking.

(d) Forward Error Correction (FEC)

(e) FEC is a capability not supported by the baseline but available for deployment.

(f) FEC will recover lost packets on a link by sending extra "parity" packets for every pre-defined group of four packets. The receiving SD-WAN router will recover any lost packet from the group, using the received parity packet and performing an XOR calculation. This delivers user experience quality that is generally in line with MPLS. The Figure below demonstrates FEC and shows application packet number three lost on the WAN link:

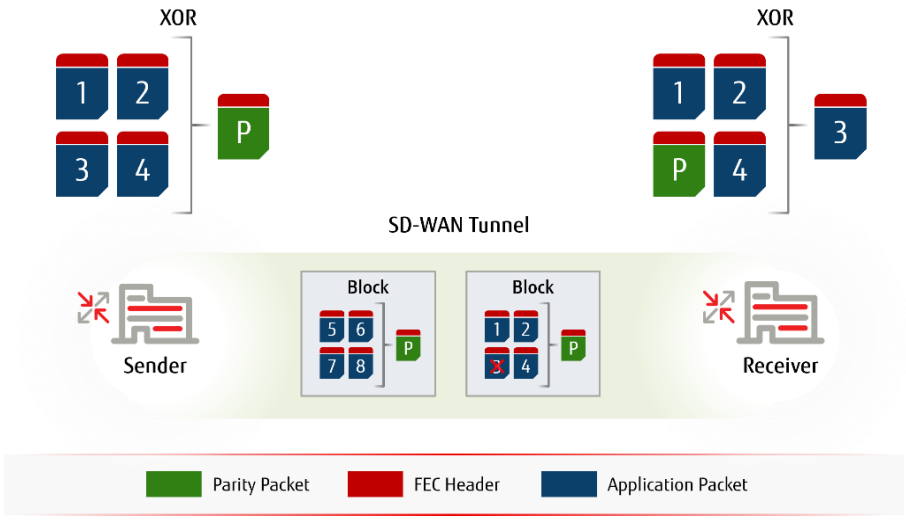

Figure 15 Forward Error Correction

(g) FEC will incur bandwidth overhead (please seek clarity from Fujitsu as to impact) as it sends an extra packet for every four end user packets, but it does not duplicate every packet.

2.6.8 WAN Interface

(a) Each SD-WAN Edge device is capable of supporting up to four WAN network connections.

(b) The SD-WAN Edge solution supports 100Mbps, 1Gbps and 10Gbps Ethernet LAN and WAN interfaces (dependent on device selection). Sub-rate network connectivity is typically used to match the site bandwidth requirements of the site, ensuring value for money.

(c) The SD-WAN router WAN interfaces can be shaped to the PSN/Internet CIR to ensure correct application of Quality of Service (QoS) and Application Quality of Experience (AppQoE) policies in the edge router.

(d) The SD-WAN network interfaces will be shaped in accordance with the information provided by the Customer

(e) The SDWAN edge device uses the following WAN interface configurations:

(i) WAN interface is 10/100/1000baseT (RJ45)

(ii) Auto-negotiation.

(iii) Option to manually set speed (100Mbps or 1Gbps) and duplex (full, half or both)

(iv) Single WAN interface per edge device.

(v) No WAN interface shaping

2.6.9 Overlay Resilience

(a) The SD-WAN edge router optimises connectivity and autonomously resolves many network availability and performance issues as they occur, improving the end user Quality of Experience (QoE) and improving productivity.

(b) Where a Customer site has multiple WAN connections, the SD-WAN overlay network will be provided over all links, including the creation of multiple paths between two sites, in accordance with the network topology templates.

(c) Multiple overlay paths are used to facilitate improvements in performance, resilience and capacity.

(d) The following Figure shows how a remote site connects to two datacentres, with resilient paths to both sites using different WAN networks.

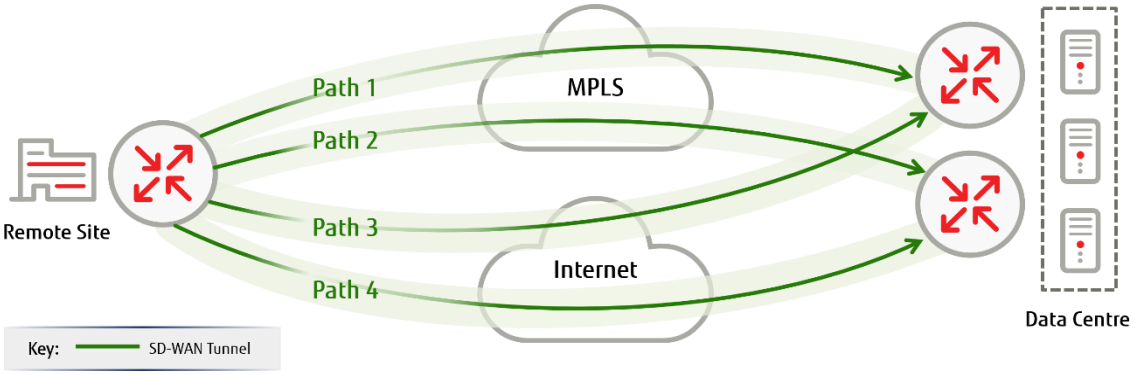

Figure 16 - SD-WAN Overlay Path Resilience

(e) Each overlay tunnel is monitored for liveliness and path quality using the Bi-directional Forwarding Detection (BFD) protocol.

(f) If BFD detects loss of connectivity, traffic will be re-distributed across the available links.

(g) By default, BFD is configured to monitor each link once every second, and will declare loss of connectivity after the loss of seven BFD messages, enabling failover in less than eight seconds.

2.6.10 Internet Breakout and Cloud Connectivity

(a) The SD WAN solution supports the capabilities for cloud connect and internet breakout as illustrated in the following figure.

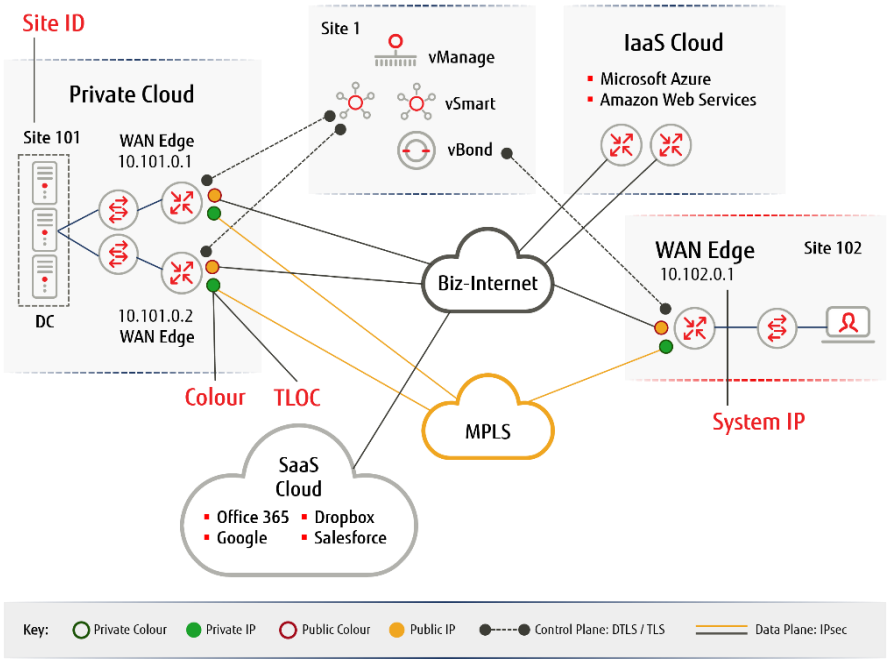

Figure 17 - SD-WAN overlay Cloud Connectivity

(b) The SD WAN solution supports the capability for native integration with third party cloud solutions including:

(i) Integration with public cloud Secure Internet Gateways.

(ii) Direction of application flows to regional colocation centres to consume Software as a Service (SaaS) and Infrastructure as a Service) applications to realise network level service chaining.

(iii) Native integration with Azure Virtual WAN, Google Cloud and Amazon Web Services (AWS).

(c) The SD WAN solution supports the capability for local, central, regional and SIG Internet breakout options, illustrated in the figures below.

(i) Local breakout provides direct internet access at the site

(ii) Remote breakout provides direct internet access at a regional hub site that is shared by spoke sites connected via the SD WAN overlay.

(iii) Central breakout provides direct internet access at a datacentre shared by remote sites connected via the SD WAN overlay.

(iv) SIG internet breakout forwards traffic to the nearest Zscaler Secure Internet Gateway Point of Presence.

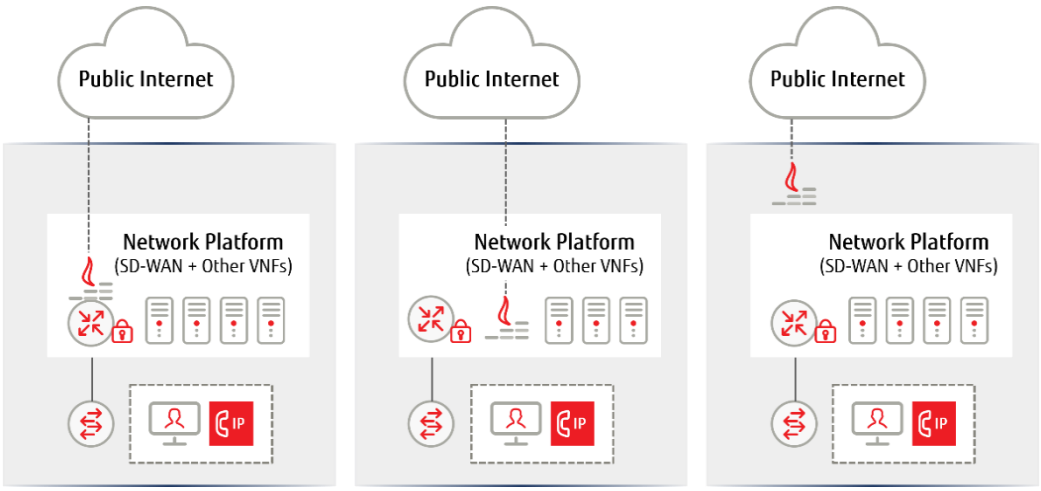

Figure 18 - Local Internet Breakout

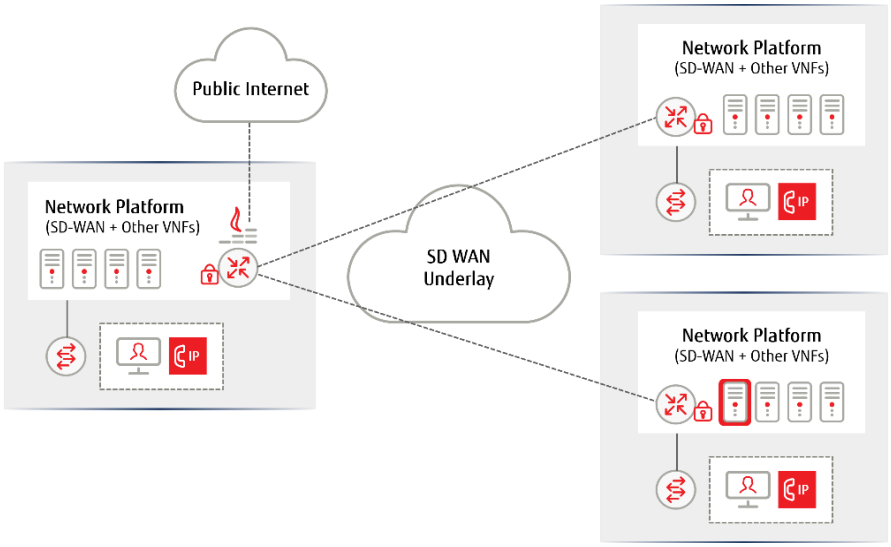

Figure 19 - Regional/Central Internet Breakout

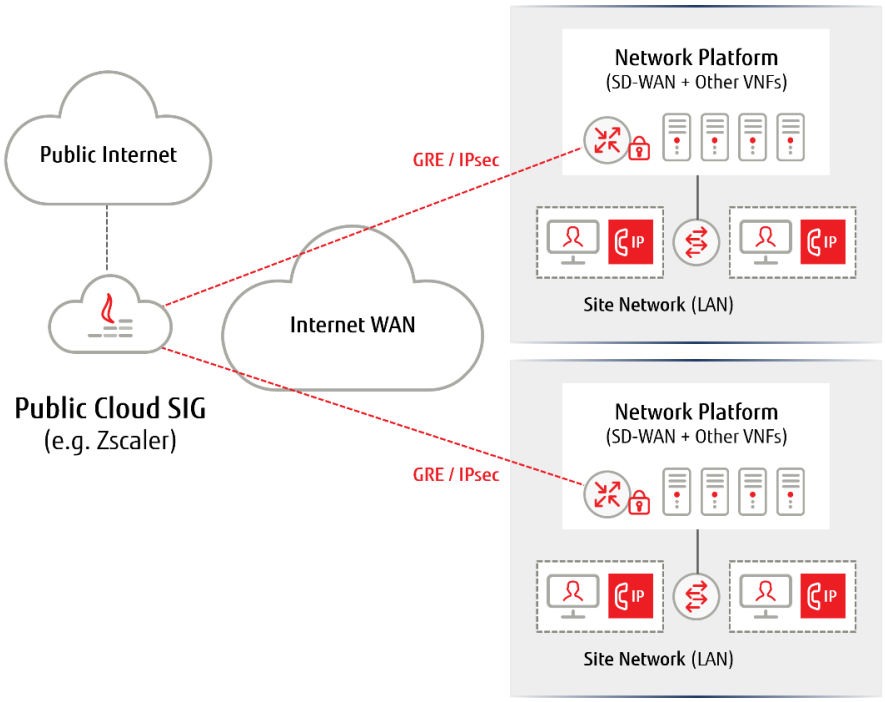

Figure 20 - SIG Breakout

(d) The SD WAN solution supports site-wide and per-VPN policies to determine the traffic and application that are allowed to use the breakout. Both static and dynamic modes are supported.

(i) With static breakout, fixed policies are defined to forward internet bound traffic to a defined breakout interface.

(ii) With dynamic breakout, the internet policy defines multiple breakout interfaces. If the SD WAN overlay BFD mechanism detects loss of connectivity or path quality on a tunnel connecting to a regional, central or SIG breakout, an alternative breakout interface is selected.

(e) For local, regional and central breakout options, a firewall should be provided at the breakout interface to protect Customer sites against internet threats.

(f) Integration of Next Generation Firewall functionality with the SD WAN and uCPE edge device is subject to change request.

(g) Support for SD WAN integration with 3rd party cloud providers using onRamp will be subject to a change request.

(i) Note: SDWAN edge devices are being provisioned in AWS in accordance using AWS infrastructure.

(h) Support for SD WAN internet and cloud breakout will be subject to a change request.

2.6.11 SD-WAN Templates and Policies

(a) The SD-WAN management solution uses configuration and policy templates to simplify and improve management process accuracy.

(b) Templates enable new sites and features to be added fast without requiring complex configurations through a command line interface (CLI).

(c) The template approach reduces the risk of mis-configuring devices through CLI typo errors at deployment or in in service. Using a template also means that configurations are inherently compliant with the solution design and security assurance requirements.

(d) Templates are applied from vManage. A vManage template can be attached to multiple WAN Edge routers simultaneously. When changes are made to configuration templates, these changes are automatically propagated to all attached SD-WAN Edge routers.

(e) There are two types of configuration templates:

(i) Feature templates help build individual components of the router configuration, such as segmentation, interfaces, system, routing, logging, and device access.

(ii) Device templates provide the framework for the entire router configuration and are made up of feature templates. Templates are flexible and allow for highly customisable router configurations. Efficient device templates design allows minimal touch configuration of thousands of devices. When making an update to a template, the changes are propagated immediately to the SD-WAN Edge routers. In case of configuration errors, the template configuration rolls back to its previous state, protecting the system against human errors.

(f) Device templates reference a series of feature templates that make up the entire configuration of a device. The device templates include the following information:

(i) Basic information and IP address Management

(ii) Transport (WAN) and management VPN

(iii) Service (Customer) VPN

(iv) System templates.

(g) Feature templates allow simple and repeatable configuration of system level features, including:

(i) System identification

(ii) Logging

(iii) Authentication, Authorisation and Accounting (AAA)

(iv) Bi-directional Forwarding Detection (BFD) monitoring

(v) Overlay Management Protocol (OMP)

(vi) Security

(vii) Archive (optional)

(viii) Network Timing Protocol (NTP)

(ix) Virtual Private Network (VPN)

(x) Border Gateway Protocol (BGP) and Pen Shortest Pat First (OSPF) routing

(xi) VPN Interface configuration

(xii) Dynamic Host Configuration Protocol (DHCP server (optional)

(xiii) Login banner

(xiv) Local policy (optional)

(xv) Simple Network Management Protocol (SNMP)

(xvi) Bridge (optional).

(h) VPN interface templates and routing protocol templates, such as BGP and OSPF, are configured under a VPN. DHCP server feature templates are configured under a VPN interface.

(i) SD-WAN policies control data traffic flow among SD-WAN Edge routers in the SD-WAN fabric. Policies relate to:

(i) Topology

(ii) Traffic flow

(iii) Local sites.

(j) Topology policies: Centralised control policies operate on the routing and Transport Locator (TLOC) information within OMP and allow customisation of routing decisions. These policies can be used in configuring traffic engineering, path affinity, service

insertion, and different types of VPN topologies (including full-mesh, hub-and-spoke or regional mesh).

(k) Traffic flow policies: Data traffic policies influence the flow of traffic through the network, based on application signatures, fields in IP headers, or which VPN segment the traffic is using. Centralised data policies are used in configuring zone based firewalls, service chaining, traffic engineering, and quality of service (QoS). They include Application-Aware Routing to apply SLAs for applications and traffic steering, while activating AppQoE features, such as packet duplication.

(l) Locally significant policies: Localised policies are used to handle traffic at a specific site. These include Access Control Lists (ACLs), Quality of Service (QoS), and route maps for OSPF, BGP or EIGRP.

(m) Policies are defined by the administrator using the policy wizard under the configuration menu of vManage. Centralised policies are applied by vManage to the vSmart controllers and localised policies are applied from vManage directly to the WAN Edge router.

2.6.12 Network Utilisation reports

(a) Fujitsu will provide a management summary of monthly network utilisation consumed per edge device. This will be composed in a single report to the Authority. Charges will be issued on the 16th day of the preceding month using the Fujitsu template reports.

(b) Utilisation reports will be created from data collected from each edge device, which will provide the inbound and outbound Network consumption measured in bytes for each polling period.

(c) The report will be made available through the LiveNX Customer Portal and will provide a graphical view of the peak and average utilization for each bin across the user selected report duration.

2.6.13 Enterprise Firewall

(a) The Enterprise Firewall is an optional service with Application Awareness and uses a flexible and easily understood zone-based model for traffic inspection, compared to the older interface-based model.

(b) A firewall policy is a type of localized security policy that allows stateful inspection of TCP, UDP, and ICMP data traffic flows.

(c) Traffic flows that originate in a given zone are allowed to proceed to another zone based on the policy between the two zones.

(d) A zone is a grouping of one or more VPNs.

(e) Grouping VPNs into zones allows security boundaries to be established in the overlay network, enabling control of all data traffic that passes between zones.

(f) Zone configuration consists of the following components:

(i) **Source zone** — A grouping of VPNs where the data traffic flows originate. A VPN can be part of only one zone.

(ii) **Destination zone** — A grouping of VPNs where the data traffic flows terminate. A VPN can be part of only one zone.

(iii) **Firewall policy** — A security policy, similar to a localized security policy, that defines the conditions that the data traffic flow from the source zone must match to allow the flow to continue to the destination zone.

(1) Firewall policies can match IP prefixes, IP ports, the protocols TCP, UDP, and ICMP, and applications.

(2) Matching flows for prefixes, ports, and protocols can be accepted or dropped, and the packet headers can be logged.

(3) Nonmatching flows are dropped by default.

(4) Matching applications are denied.

(iv) **Zone pair** — A container that associates a source zone with a destination zone and that applies a firewall policy to the traffic that flows between the two zones.

(g) Matching flows that are accepted can be processed in two different ways:

(i) **Inspect** — The packet's header can be inspected to determine its source address and port. When a session is inspected, it is not necessary to create a service-policy that matches the return traffic.

(ii) **Pass** — Allow the packet to pass to the destination zone without inspecting the packet's header at all. When a flow is passed, no sessions are created. For such a flow, a service-policy must be created that will match and pass the return traffic.

2.7 Application Monitoring

2.7.1 Overview

(a) The SD WAN service includes an analytics service to support monitoring and analysis of Customer application flows.

(b) The analytics service provides the Customer access to the LiveNX network and application performance monitoring platform.

(c) Fujitsu use the LiveNX platform for monitoring of the SD-WAN service:

(i) uCPE availability monitoring and SLA reporting

(ii) SD-WAN overlay monitoring and SLA reporting

(iii) SD-WAN edge utilisation

(d) The Customer shall be provided access to the LiveNX platform to support the following capabilities:

(i) Correlation of multiple data sets to provide views, graphs and maps to illustrate the current state of applications and network performance.

(ii) Application visibility and troubleshooting to gain a deep understanding of application traffic with full visibility of protocol and application type including video, voice, instant messaging, file transfer, etc.

(iii) Application analysis to trouble shooting.

(iv) Analysis of how the SD-WAN network is being used, how applications are performing, and which sanctioned or unsanctioned applications are being used.

2.7.2 Management Interface

(a) LiveNX is managed via the LiveNX JAVA client or web User Interface (UI).

(b) LiveNX shall be integrated with the SD-WAN vManage appliance through the REST API management interface.

2.7.3 Network Data

(a) LiveNX shall collect the following SD-WAN performance data:

(i) uCPE availability

(ii) SD-WAN overlay BFD performance metrics (loss, latency, jitter and path failure)

(iii) SD-WAN edge router WAN utilisation statistics

2.7.4 Application Flow Data

(a) Application flow data shall be collected directly from SD-WAN edge devices using cflowd.

(b) Cflowd is a flow analysis tool, used for analysing NetFlow traffic data. It monitors traffic flowing through the SD-WAN C8000v edge devices in the overlay network and exports flow information to a collector, where it can be processed by an IP Flow Information Export (IPFIX) analyser. For a traffic flow, cflowd periodically sends template reports to the flow collector. These reports contain information about the flows and the data is extracted from the payload of these reports.

(c) Cflowd traffic flow monitoring is equivalent to Flexible Netflow (FNF). The cflowd software implements cflowd version 10, as specified in RFC 7011 and RFC 7012. Cflowd version 10 is also called the IP Flow Information Export (IPFIX) protocol.

(d) The SD-WAN edge router cflowd-template defines the location of cflowd collectors, how often sets of sampled flows are sent to the collectors, and how often the template is sent to the collectors (on Cisco vSmart Controllers and on Cisco vManage).

2.8 Customer Portal Access

(a) The SD WAN service shall provide the following Customer Portal access to user authorised by the LEC Authority:

(i) vManage

(ii) LiveNX Client

(b) Access shall be authenticated and authorised in accordance with the Role Based Access Control (RBAC) process, which includes:

(i) All accounts shall be assigned to an individual person.

(ii) Access shall be provided on a least privilege basis.

(c) Access shall be provided to vManage and LiveNX Web server interfaces via a web server proxy, which will be accessed from End User Devices via the SD WAN overlay (service VPN).

(d) The Customer Portal shall provide a Remote Desktop facility to support LiveNX access using the Java client.

(e) End user devices shall be authenticated by policing certificates.

(i) Customer to provide policing root CA certificates

2.9 PSN and Internet Underlay

2.9.1 Underlay Network

(a) The SD-WAN uses standards-based IP/Ethernet WAN interfaces with Network Address Translation to enable use of any IP-based networking technology, including PSN and MPLS Virtual Private Networks (VPNs) using private IP address space, and the Internet using public IP addressing.

(b) The SD-WAN shall use the following underlay networks provided by the Customer:

(i) Public Sector Network

(ii) Internet ISP

(iii) 5g ISP Connectivity

(iv) MPLS networks

(c) The SDWAN Customer shall provide a WAN switch to facilitate connectivity of SDWAN edge devices to the PSN CPE in both standard (active and warm standby) and HA (active / active) modes of operation.

(d) Each site shall be operated in one of the following modes of operation:

(i) Standard, non-resilient active uCPE plus an active warm standby device with a single PSN/internet network connection.

(ii) High Availability, resilient active/active uCPE with resilient PSN/internet network connections and resilient PSN CPE operating in HSRP mode.

(e) Support for other operating modes, including the following, is subject to change control:

(i) High Availability, resilient uCPE with single PSN/internet network connections

(ii) High availability sites using TLOC extensions.

(f) The level of resilience provided by the network connectivity (e.g. diverse or fully diverse) shall be determined by the PSN service provided by the site and defined in the MSL.

(g) The Customer shall provide the following network connectivity information for each site:

(i) IP subnet, mask and address allocations

(ii) Circuit type (Ethernet)

(iii) Link speed (10Mbps, 100Mbps, 1Gbps, 10Gbps)

(iv) Committed rate (<= link speed)

(v) Physical coding sublayer (10/100/1000baseT, 1000baseX, 1000baseT, 10GbaseX)

(vi) Physical medium (copper, single mode fibre, multi-mode fibre)

(vii) Physical layer (SR, LR)

(h) The SD-WAN solution supports Network Address Translation at public internet interfaces.

2.9.2 Network Utilisation

(a) The SD-WAN solution operates by default with ECMP for WAN connectivity, meaning it will load-balance two or more WAN circuits, thereby providing active-active load balancing by default. If a Customer wishes to ensure traffic is routed via one particular transport network, there are preference mechanisms that force traffic to use one link rather than another.

(b) The SD-WAN solution fully supports load balancing per flow within the overlay network. The SD-WAN fabric supports numerous methods of mapping any given application onto any of the available WAN transport links in any of the required service (LAN) side VPN:

(i) Per-session active-active load sharing across multiple transports irrespective of the transport type (PSN or Internet).

(ii) Per-session active-active weighted load sharing across multiple transports where certain configured ratios are applied for proportional traffic distribution.

(iii) Active/standby for pinning application traffic to specific transport links with or without failover.

(iv) SLA-based application aware routing, where the choice of transport is governed by meeting (or not meeting) desired loss, latency and jitter characteristics for the given application.

(c) These forwarding methods are not mutually exclusive and can be leveraged all at once, while acting on different types of application traffic as defined by 6-tuple matching (including DSCP value) or DPI signatures.

2.10 Customer LAN

2.10.1 Customer LAN Integration

2.10.2 The SD-WAN supports features that simplify integration with Customer Local Area Networks (LAN) and applications.

(a) Use of standards-based IP/Ethernet interfaces, OSPF/BGP routing protocols and the Virtual Router Redundancy Protocol (VRRP) allow connectivity with standard LAN topologies.

(b) The Customer LAN shall provide any layer 2 connectivity required to support VRRP in resilient uCPE configurations.

(c) Virtual Routing Functions maintain the logical separation of Customer Closed User Group (CUG) traffic.

(d) Layer 2-4 classification enables application-based policies for encrypted and unencrypted applications flows, while Deep Packet Inspection supports further analysis of unencrypted application flows.

(e) The Customer shall provide the LAN integration requirements to support discovery activities of each site Low Level Design, which shall include the following:

(i) Customer VLANs and subnets

(ii) Customer LAN gateway IP address allocation

(iii) Static routes

(iv) Dynamic routing details

(v) VRRP requirements

2.10.3 Customer VPNs

(a) The SD WAN solution allows Customer traffic to be mapped to different VPN segments across the secure fabric using a single set of IPsec tunnels among the sites. VPN segments provide strict logical separation of traffic, have distinct VPN

topologies, overlapping IP addresses and unique application and security policies. VPN segments will be connected with routing and security policies at individual sites.

(b) The LEC SD-WAN solution implements segmentation at the edge and allows for scalability into 100s of VPNs.

(c) By default the SD WAN shall provide separate VRFs for the following Customer LAN segmentation:

(i) Customer defined service VPN LANs/VLANs

2.10.4 Network Address Translation (NAT)

(a) The SD WAN solution supports Network Address Translation (NAT) on the Customer LAN side. This will allow the support of overlapping IP subnets and re-use of existing IP subnets and schema on the Customer LAN.

(b) The LAN side addresses are redistributed to the overlay and advertised to all the remote branches using the Overlay Management Protocol (OMP). Thus, the remote host is aware of the path to reach inside hosts.

(c) NAT will ensure that IP addresses in the overlay transport VPN are unique.

2.10.5 DHCP

(a) The SD-WAN edge device supports the capability to provide DHCP services for the Customer LAN.

(b) No DHCP services are enabled by default for the SDWAN service.

2.11 AWS Edge

(a) SDWAN will provide support for edge devices deployed in the AWS cloud infrastructure as illustrated in the following figure.

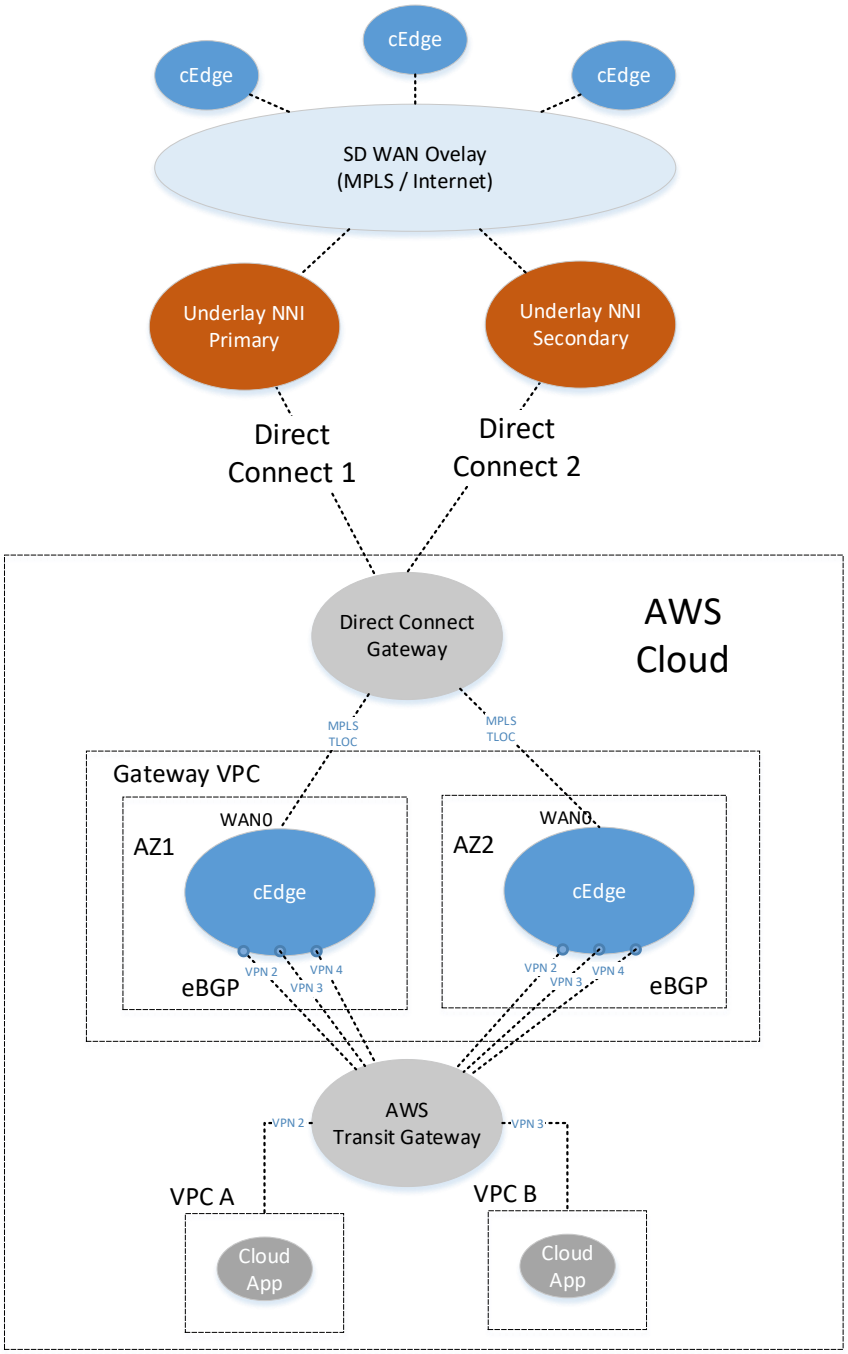

Figure 21 SD WAN – AWS TGW Integration

(b) The Customer is responsible for providing underlay connectivity between the SDWAN and AWS infrastructure.

(i) AWS is connected to the PSN via two AWS Direct Connect services providing 1:1 (active/active) resilience.

(c) The SDWAN service will provide the following information to enable initial build of the AWS edge instance:

(i) SDWAN DNA licensing

(ii) Edge device UUID

(iii) Edge device OTP

(iv) Entrust root CA trust chain

(d) The SDWAN service will be responsible for on-boarding and managing the SDWAN cEdge instance.

(e) The SDWAN edge will establish eBGP peering with the Transit Gateway and advertise learned routes into OMP.

2.12 Availability & Resilience

2.12.1 Overview

(a) The SD-WAN service provides a managed overlay service and provides availability service level agreements for the following aspects:

(i) SDWAN Core Platform availability, measured as the ability to maintain SDWAN Overlay connectivity

(ii) SD-WAN edge device availability

(b) The following aspects of the service are provided by the SDWAN Customer and are outside the scope of the service:

(i) Power and accommodation of SD-WAN edge devices at Customer premises, including:

(1) Building integrity

(2) Heating, ventilation and air-conditioning

(3) Power supply

(ii) Underlay network connectivity

(iii) WAN and LAN cabling and switching infrastructure

(iv) LAN switching

(v) Onsite assistance (excluding break fix)

(c) The SD-WAN solution supports Core Platform, MANO, edge device, and overlay resilience as illustrated in figure 22 SD-WAN Redundancy Mechanisms (note diagram redacted).

Figure 22 - SD-WAN Redundancy Mechanisms

(a) The SDWAN will inherit transport layer resilience provided by the PSN/internet underlay, which is supplied by the SDWAN Customer.

(b) The SDWAN provides fully meshed overlay connectivity between all sites in each service VPN. This enables application providers to provide dual homed connectivity to end user sites. The mechanism for switching end users from one application datacentre to another is outside the scope of the SDWAN service.

(c) At Customer sites, SD-WAN edge devices shall be deployed in one of the following modes:

(i) High Availability (HA) mode, with two devices providing active/active operation

(ii) Standard Availability mode with one active device and one warm standby device.

2.12.2 Edge High Availability

(a) High Availability mode can be used at any site.

(b) High Availability deployments protect against multiple layers of failure including:

(i) Loss of power

(ii) Device failure

(iii) WAN link failure

(iv) PSN / MPLS and Internet underlay network failure

(v) SD-WAN overlay tunnel failure.

(c) The provision of separate power supplies is the responsibility of the Customer site.

(d) The remote edge (VEP-4600) provides dual power supply modules and supports full operation with one operational power supply.

(e) In the event of a dual power or device failure Customer LAN traffic will be routed via the remaining device and the WAN connections supported by that device.

(f) Resilience against underlay failure is the responsibility of the PSN/MPLS/internet service provider. For HA sites with dual network circuits and CPE, resilience will be provided from Customer edge to Customer edge.

(g) The High Availability solution is based on the following principles.

(i) ***Device redundancy***: In High Availability mode, a primary and secondary device are provided, which operate in active/active mode. The active/active mode of operation permits usage of the resilient PSN connectivity using HSRP protected CPE and WAN Ethernet switching infrastructure.

(ii) The connectivity of traffic from both devices in a HA configuration to all available WAN networks is illustrated in the following figure 23 HA device resilience (redated) and figure 24 (redacted)

Figure 23 – HA Device Resilience (SDWAN Overlay)

Figure 24 - HA Device Resilience (Pass-through)

(i) Active/active mode of operation requires SDWAN and EC hypervisor licences for each edge device.

(ii) The HA mode of operation has been defined by the HO LEC technical team.

(iii) ***Edge - Service VPN VRRP***. The HA mode of operation requires the edge devices to use VRRP in each service VPN. VRRP provides a VIP as the SDWAN default gateway for Customer LAN devices (e.g. forces firewall), which can be reached from the Customer LAN environment via an external LAN switch.

(iv) The HA pair uses VRRP WAN/LAN prefix tracking to monitor the status of the edge device WAN and LAN connectivity. If the WAN/LAN connectivity goes down, VRRP will move the VIP to direct Customer LAN traffic to the other device and route traffic via the alternative WAN/LAN link.

(v) ***PSN CEP – HSRP*** SDWAN edge devices will forward traffic to the CPE Virtual IP (VIP) address in accordance with the ARP table. In the event of a CPE device failure, the SDWAN ARP table will be updated in accordance with gratuitous ARPs sent by the active PSN CPE.

(vi) ***Forces Firewall – VRRP:*** SDWAN edge devices will forward traffic to the Forces Firewall Virtual IP (VIP) address in accordance with the ARP table. In the event of a firewall device failure, the SDWAN ARP table will be updated in accordance with gratuitous ARPs sent by the active firewall.

(vii) ***SDWAN OMP routing***. The SDWAN OMP routing protocol ensures rapid recovery from both direct and indirect failure. To provide a resilient control plane, the solution regularly monitors the status of all WAN Edge routers in the network and automatically adjusts to changes in the topology as routers join and leave the network.

2.12.3 Edge Standard Availability

(a) Standard Availability mode can be used at any site, and will be used as defined by the CA MSL. Standard deployments have less resilience against failures than HA deployments, but do maintain protection against multiple layers of network failure including:

(i) Power supply failure.

(ii) WAN link failure (if the site has multiple links)

(iii) PSN / MPLS and Internet underlay network failure

(iv) SD-WAN overlay tunnel failure.

(b) The provision of separate power supplies is the responsibility of the Customer site.

(c) The remote edge (VEP-4600) provides dual power supply modules and supports full operation with one operational power supply.

(d) The SD-WAN solution supports.

(i) ***Primary device***: In Standard mode, a single primary device provides full operational service.

(ii) ***Warm standby device***: In Standard mode, a secondary standby device is actively managed, but not connected to the Customer LAN infrastructure. The warm standby provides an on-site break fix device that is ready to replace a failed primary device. The warm standby device is fully managed and maintained with the latest patches and software releases, monitored for defects and security incidents, and synchronised to the Customer datacentre clock. In the event of a primary device failure, the Customer LAN infrastructure needs to be moved to the secondary device. The DJSC will remotely copy the failed primary device configuration to the warm standby device. The Standard mode of operation permits usage of a single CPE interface using WAN Ethernet switching infrastructure (provided by the SDWAN Customer). The connectivity of traffic from both devices in a Standard configuration is illustrated in the following figures 25 standby device connectivity (redacted) and 26 standard device resilience.

Figure 25 – Standard Device Connectivity (SDWAN overlay)

Figure 26 – Standard Device Resilience (Pass-through)

(iii) Standard mode of operation requires SDWAN and EC hypervisor licences for each edge device.

(iv) The Standard mode of operation has been defined by the HO LEC technical team.

(e) In order to meet service restoration targets of less than 2 hours, local support or Customer Smart Hands are required to move Customer LAN connections from a failed primary edge device to the warm standby device.

2.12.4 Core Platform - SDWAN MANO Availability

(a) The SDWAN provides a high availability management and orchestration (MANO) platform located in two separate datacentres (DC1 and DC2) using the following resilience capabilities:

(i) Active/active orchestration (vBond), control plane (vSmart) and DNS.

(ii) Active/standby management (vManage)

(iii) Resilient PSN connectivity to each SDWAN MANO

(iv) Dynamic SDWAN routing (OMP)

(v) Autonomous edge and datapath operation in the event of MANO failure (graceful restart).

(b) Once the SD-WAN edge router has established the SD WAN overlay, in accordance with centralised templates and policies, each edge router uses local packet processing and forwarding to route traffic between Customer sites. In the event of

control plane failures the edge devices operate in Graceful Restart mode, which caches local routing and security parameters, allowing the device to continue to operate in the data plane.

(c) ***vBond***: The vBond orchestrator operates in active/active mode, with one vBond instance in each datacentre. When establishing management and control plane connectivity, edge devices requests primary and secondary vBond IP address details from the dedicated SDWAN DNS server.

(d) ***SDWAN DNS***: An active DNS server is provided in each datacentre. If an edge device does not receive vBond details from the primary DNS server, it will send a request to the secondary server.

(e) ***vSmart***: The vSmart controller operates in active/active mode, with one vSmart instance in each datacentre. When establishing management and control plane connectivity, edge devices establish an active connection with both vSmart controllers. If one of the vSmart controllers fails, the other one seamlessly take over handling control of the network.

(f) For correct operation, the control policies in each vSmart controller must be identical. To remain synchronized with each other, the vSmart controllers establish a full mesh of DTLS control connections, as well as a full mesh of OMP sessions, between themselves. Over the OMP sessions, the vSmart controllers advertise routes, TLOCs, services, policies, and encryption keys. It is this exchange of information that allows the vSmart controllers to remain synchronized.

(g) ***vManage***: vManage operates in active/standby mode, with one vManage instance in each datacentre. In normal operation only the active vManage is connected to the transport network and vBond established edge connectivity with this vManage instance. In the event of a vManage failure, a DJSC administrator manually copies the latest configuration database to the standby vManage and performs a Disaster Recovery failover to establish the standby vManage as the active controller. In order to support this process, the SDWAN toolset server takes regular backups of the configuration database (at 8 hour intervals) and copies it to the opposite datacentre.

(h) ***SDWAN OMP routing***: The SDWAN OMP routing protocol ensures rapid recovery from both direct and indirect failure. To provide a resilient control plane, the solution regularly monitors the status of all WAN Edge routers in the network and automatically adjusts to changes in the topology as routers join and leave the network.

(i) ***PSN connectivity***: The SDWAN datacentres uses BGP dynamic routing via the DC interconnect to provide resilient PSN connectivity from each datacentre, as illustrated in the following figure.

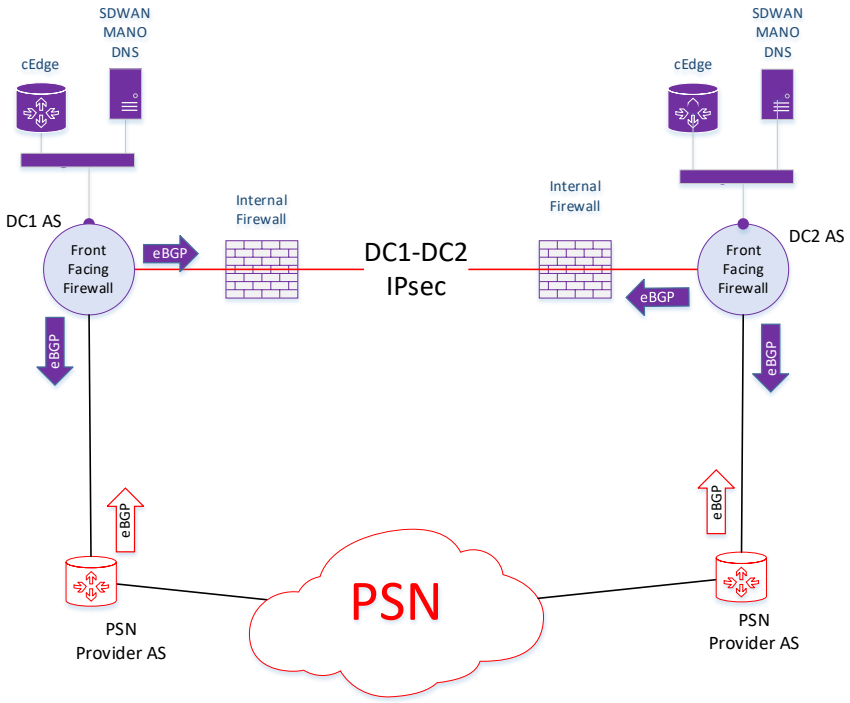

Figure 27 Connectivity Resilience - BGP Dynamic Routing

2.12.5 Core Platform – Infrastructure and Enterprise Management Availability

(a) The SDWAN Core Platform provides high availability infrastructure and Enterprise Management across two separate datacentres (DC1 and DC2) using the following resilience capabilities:

(i) Clustered server infrastructure

(ii) Primary and backup storage

(iii) Clustered and virtual chassis DC network equipment

(iv) Backup and recovery with offsite backups stored as virtual tapes in opposite datacentre

(v) High availability Enterprise Management applications

(b) The SDWAN Core Platform (GSA Customer pod) is supported by a 2+1 vSphere High Availability cluster.

(c) The SDWAN Core Platform (GSA Customer pod) provide primary and backup storage via iSCSI.

(d) All DC firewalls are operated with dual power supplies and configured in a 1+1 active/standby cluster.

(e) All DC switches are operated with dual power supplies and configured in a 1:1 active/active virtual chassis.

(f) Remote access between BRA07/BSN01 resolver groups and the Core Platform is provided via resilient network connectivity, terminal servers and Remote Desktop Services.

(g) The LiveNX performance management application is operated as a single server, with Veeam backup and replication providing server resilience and datacentre DR.

(h) The Elasticsearch Enterprise Management toolsets is configured as a 3 node cluster in each datacentre with cross-cluster data replication.

(i) The Zabbix Enterprise Management toolset is configured as an active/standby pair with automated failover.

2.13 Service Interruption

(a) The following table defines the service impact as a result of a single network, device or controller component failure in the SDWAN service.

LEC SDWAN Component	LEC SDWAN Service Impact (Single Component Failure)	Resilience Mechanism
SDWAN Edge (Standard)	<4 hours	Warm standby
SDWAN Edge (HA)	<5 minutes	VRRP
Single Edge power supply	None	VEP-4600 dual power supplies
Customer site PSN circuit (standard)	Refer to PSN provider SLA	Refer to PSN provider solution
PSN network	Refer to PSN provider SLA	Refer to PSN provider solution
SDWAN MANO – vBond	None	Active / active vBond
SDWAN MANO – vSmart	None	Active / active vSmart
SDWAN MANO – vManage	Loss of SDWAN management No loss of SDWAN Overlay data path	vManage DR failover
SDWAN DNS	None	Active / active DNS
LiveNX	Loss of application monitoring No loss of SDWAN Overlay data path	LiveNX VM replication
ISE	None	Active / active ISE
NTP	None	SRX cluster NTP server active (DC1) / active (DC2)
Zabbix (SNMP)	None	Zabbix active /standby automated failover

LEC SDWAN Component	LEC SDWAN Service Impact (Single Component Failure)	Resilience Mechanism
Elasticsearch (Syslog)	None	3 node cluster

Table 2 –SDWAN Component Failure - Service Impact

(b) Typical service interruption times due to common failure modes:

(i) Failover due to device reboot or power failure on a HA platform will be in the order of < 5 minutes, which is defined by the time for VRRP to detect the failure and cause a routing change.

(ii) An SD-WAN router reboot (as an example, for software upgrade) will typically be 8 minutes.

(iii) The uCPE platform will automatically reboot after restoration of power, resulting in a reboot of typically 8 minutes.

(iv) Failover due to interface failures will be in the order of <5 minutes, which is defined by the time for VRRP to detect the failure and cause a routing change.

2.14 Backup and Recovery

2.14.1 VM Backup and Recovery

(a) The core Platform provides a Veeam backup and replication solution for Virtual Machines and specified files.

(b) The SDWAN Core Platform VMs are automatically backed up daily by the Veeam backup solution.

(c) Backups are stored locally and offsite in the opposite datacentre.

(d) Backups are stored as encrypted virtual tapes in dedicated storage partitions.

2.14.2 MANO Configuration

(a) In addition to VM backups, the LEC SDWAN vManage configuration database is backed up at 8 hour intervals.

(b) SDWAN configuration database backups are performed automatically by the toolset server, stored locally and copied to the opposite datacentre.

(c) In the event of a vManage DR failover, the latest vManage configuration database is restored on the standby vManage.

2.15 Disaster Recovery

(a) The SD-WAN service offers a Business Continuity and Disaster Recovery (BCDR) with the following approaches:

(i) DR datacentre

(ii) Standby resolver group locations

(iii) VM backup and replication

(iv) High Availability Core Platform

(b) The Core Platform is deployed in geographically separate primary and secondary datacentres in [redacted]

(c) The Core Platform primary resolver group is located in [redacted], with DR resolver group facilities provided in [redacted]

(d) The Core Platform backup and replication solution allows recovery of individual VMs in both datacentres from backups stored in both datacentres.

(e) The SDWAN Core Platform also provides high availability resilience without invoking DR processes such as VM recovery:

(i) Clustered server infrastructure

(ii) Clustered and virtual chassis DC network equipment

(iii) High availability MANO and Enterprise Management toolsets

(f) Dependent on the nature of the functionality provided, individual toolsets are protected in one of the following active/active or active/standby modes of operation:

(i) Active/active operation

(ii) Active/standby with automated stateful database synchronisation or replication and administrator failover

(iii) Active/standby with automated stateful database backup and administrator failover

(iv) Active/- with automated virtual machine backup and manual recovery

(v) Active/standby with automated failover

(g) The following table summarises the High Availability and Disaster Recovery approaches for individual management functions in the SD-WAN toolset between datacentres.

Toolset	Application Level DR (DC1 to DC2)	Disaster Recovery (DC2)
Remote Desktop Service	Active / Active	VM recovery
LOSS/ESO (Not currently used in LEC SDWAN)	Active/- with automated VM backup and manual recovery	VM recovery
vManage	Active / standby vManage (Automated database backup and admin failover)	VM recovery
vBond	Active / Active	VM recovery
vSmart	Active / Active	VM recovery
NTP server	Active / Active	VM recovery
Elasticsearch (Syslog)	Active / Standby (Data replication and automated failover)	VM recovery
LiveNX (Flexible Netflow)	Active/- with automated VM backup and manual recovery	VM recovery
ISE (RADIUS/ TACACS+)	Active / Active	VM recovery
Zabbix (SNMP traps)	Active / Standby (Automated database backup and admin failover)	VM recovery

Table 3 – SDWAN Core Platform Toolset Disaster Recovery Approach

2.16 Performance Management

2.16.1 Current Growth and Expected Growth

(a) The SD-WAN edge routers are deployed using the Virtual Edge uCPE approach.

2.16.2 Scalability

(a) The SD-WAN deployment is based on multiple instances of SD-WAN management and controller appliances configured in high availability mode, with individual appliances being able to support up to 2000 edge devices.

(b) All uCPE deployment support 8x vCPU, 16GB RAM and 240GB storage with a C8000v virtual router, which supports a throughput of up to 2Gbps.

2.17 Security

2.17.1 Control Plane Security

(a) The LEC SD-WAN control plane uses digital certificates with 2048-bit RSA keys to authenticate the SD-WAN edge routers in the network. The digital certificates are created, managed, and exchanged by standard components of the public key infrastructure (PKI):

(i) Public keys — These keys are generally known.

(ii) Private keys — These keys are private. They reside on each SD-WAN router and cannot be retrieved from the router.

(b) Certificates are signed by a root certification authority (CA). The trust chain associated with the root CA needs to be present on all Cisco SD-WAN controllers and routers.

(c) For the SDWAN, the root CA is provided by Fujitsu's CA service.

(d) For vSmart controllers, vBond orchestrators, vManage systems, the certificates are managed manually. The Cisco SD-WAN software generates a unique private key–public key pair for each software image. The network administrator requests a Certificate Signing Request for each controller, which is sent to the issuing sub-CA of the root CA trust chain for signing and manually installed on the corresponding controller virtual appliance.

(e) Control plane encryption is done by either DTLS, which is based on the TLS protocol, or TLS. These protocols encrypt the control plane traffic that is sent across the connections between SD-WAN devices to validate the integrity of the data. TLS uses asymmetric cryptography for authenticating key exchange, symmetric encryption for confidentiality, and message authentication codes for message integrity.

(f) For the SDWAN, TLS is used for vManage and vSmart communications and DTLS is used for vBond communications.

(g) The SD-WAN design implements control plane integrity by combining two security elements: SHA-2 message digests, and public and private keys.

(h) SHA-2 are cryptographic hash functions that generate message digests for each packet sent over a control plane connection. SHA-2 is a family that consists of six hash functions with digests that are 224, 256, 384, or 512 bits. The receiver then generates a digest for the packet, and if the two match, the packet is accepted as valid. SHA-2 allows verification that the packet's contents have not been tampered with.

(i) The second component of control plane integrity is the use of public and private keys. When a control plane connection is being established, a local SD-WAN device sends a challenge to a remote device. The remote device encrypts the challenge by signing it with its private key, and returns the signed challenge to the local device. The local device then uses the remote device's public key to verify that the received challenge matches the sent challenge.

(j) Then, once a control plane connection is up, keys are used to ensure that packets have been sent by a trusted host and were not inserted midstream by an untrusted source. The authenticity of each packet is verified through encryption and decryption with symmetric keys that were exchanged during the process of establishing the control connection.

2.17.2 Data Plane Security

(a) The underlying foundation for security in the SD-WAN data plane is the security of the control plane.

(b) Because the control plane is secure (all devices are validated, and control traffic is encrypted and cannot be tampered with) the SD-WAN devices can be confident in using routes and other information learned from the control plane to create and maintain secure data paths throughout a network of routers.

(c) The data plane provides the infrastructure for sending data traffic among the routers in the SD-WAN overlay network. Data plane traffic travels within secure Internet Security (IPsec) connections. The SD-WAN data plane implements the key security components of authentication, encryption, and integrity in the following ways:

(d) Authentication, the SD-WAN control plane contributes the underlying infrastructure for data plane security. In addition, authentication is enforced by two other mechanisms:

(i) As standard the 'traditional' SD-WAN key exchange model, the vSmarts sends IPsec encryption keys to each edge device.

(ii) In the optional pairwise keys model, the vSmart sends Diffie-Hellman public values to the edge devices and they generate pairwise IPsec encryption keys using ECDH and a P-384 curve

(iii) The LEC SDWAN uses the optional pairwise keying model.

(iv) By default IPsec tunnel connections use a modified version of the Encapsulating Security Payload (ESP) protocol for authentication on IPsec tunnels. This version of the protocol also checks the outer IP and UDP headers. Hence, this option supports an integrity check of the packet similar to the Authentication Header (AH) protocol.

(v) ***Encryption*** — Data encryption is done using the [redacted].

(vi) ***Integrity*** — To guarantee that data traffic is transmitted across the network without being tampered with, the data plane implements several mechanisms from the IPsec security protocol suite:

(1) The modified version of ESP uses an AH-like mechanism to check the integrity of the outer IP and UDP headers. You can configure the integrity methods supported on each router, and this information is exchanged in the router's TLOC properties. If two peers advertise different authentication types, they negotiate the type to use, choosing the strongest method.

(2) The anti-replay scheme protects against attacks in which an attacker duplicates encrypted packets.

2.17.3 Cipher Suites

(a) vManage, vBond, vSmart and C8000v use the cipher suites listed in the following Tables.

Cryptographic Operations	Control Plane (vBond/vSmart to C8000v)	Management Plane (vManage to C8000v)	Data Plane SD-WAN (C8000v to C8000v)	Data Plane Pairwise (C8000v to C8000v / 3rd Party)
Secure Protocols	[redacted]		[redacted]	IPsec
Digital Signature Generation and Verification	[redacted]		Control Plane	Control Plane
Key Agreement	[redacted]		vSmart control	[redacted]
Symmetric Encryption and Decryption	[redacted]		[redacted]	[redacted]
Message Authentication	[redacted]		[redacted]	[redacted]

Table 4 - Cryptographic Operations of Cisco SD-WAN

2.17.4 Internet Connectivity

(a) For internet connections used as the underlay the SD-WAN will use [redacted] level 1 compliant SD WAN encryption algorithms.

2.17.5 TLS

(a) From Service Commencement the Authority is responsible for certification and provision of TLS software meeting NPIRMT standards for the applications. The Fujitsu Managed Cisco SD-WAN will provide a connection (and routing) to support the TLS [redacted] applications in accordance with the features and functionality of the Cisco SD-WAN version software guide.

2.17.6 Edge Device Security Boundary

(a) The following figure defines the security boundary of the SDWAN edge device.

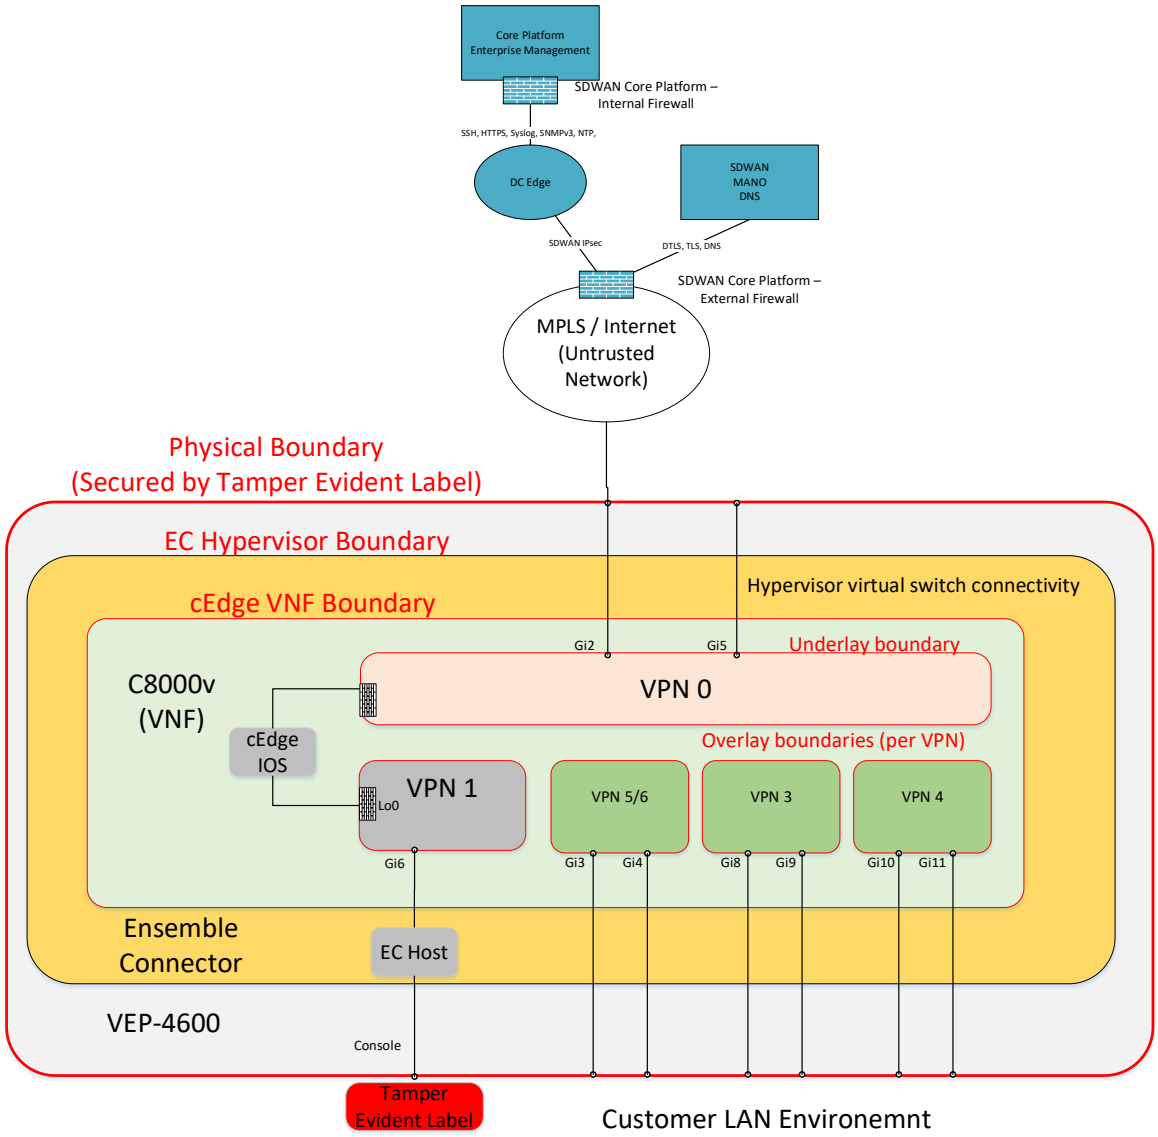

Figure 28 LEC SDWAN Edge – Port Connectivity

(b) Unused ports are not allocated to a VPN and are disabled.

(c) The only management and control plane access from the PSN/internet is the C8000v edge through DTLS, TLS, DNS.

(d) The only data plane access from the PSN/internet is the SDWAN overlay (IPsec).

(e) The SDWAN service VPNs cannot be accessed from the PSN/internet.

(f) The Adva EC host cannot be accessed from the PSN/internet.

(g) The Core Platform Enterprise Management infrastructure and toolsets can only be reached via the SDWAN overlay Telemetry VPN (VPN1).

(h) SDWAN VPN1 has no external interfaces at the SDWAN edge.

(i) The SDWAN edge includes the following firewalls on external facing interfaces:

(i) cEdge VPN0 host based firewall (PSN/internet underlay) - allow TLS, DTLS, DNS services only

(ii) cEdge VPN0 self-zone firewalls (PSN/internet underlay) - allow TLS/DTLS and DNS connectivity to/from LEC Core Platform only

(iii) cEdge vPN1 self-zone firewall (Core Platform) - allow management connectivity to/from LEC Core Platform only

(j) The Core Platform includes the following firewalls on remote edge facing interfaces:

(i) PSN external firewall - allow TLS/DTLS and DNS connectivity to/from SDWAN edge only

(ii) Core Platform internal firewall (VPN1) - allow management connectivity to/from LEC SDWAN edge only

2.17.7 Core Platform Security Boundary

(a) The following figure 29 redacted defines the security boundary of the SDWAN edge device.

Figure 29 - Core Platform Security Boundaries

(a) SDWAN MANO and DNS is located in an SDWAN demilitarized zone:

(i) Physically separate external firewall provides access to the PSN

(ii) Front facing firewall permits access to VPN0 interfaces of MANO (vManage, vBond, vSmart) and cEdge in the datacentre.

(iii) Internal firewall permits access to management VPN (VPN512) of the SDWAN MANO and DNS components; and service VPNs of the cEdge.

(b) Front facing firewall provided PRIME IPsec connectivity to front facing firewall in opposite datacentre via DC interconnect for PSN connectivity resilience.

(c) Front facing firewall only permits IPsec and IKE to be routed via internal firewall.

(d) cEdge provides controlled access to the Core Platform via SDWAN overlay service VPNs:

(i) ***VPN1*** Enterprise Management toolset

(ii) ***Customer defined service VPNs***

(e) Internal firewall permits access from SDWAN overlay (VPN1) to Enterprise Management infrastructure and toolsets:

(i) NTP,

(ii) LiveNX (cflowd, SNMPv3)

(iii) Elasticsearch (syslog),

(iv) Zabbix (SNMP),

(v) RDS (SSH),

(vi) ELS (HTTPS),

(vii) ISE (TACACS+)).

(f) The Core Platform internal firewall cannot be accessed from the PSN/internet.

3. Service Delivery, Service Packs 1 & 2

Service Management

Fujitsu is an ITIL® aligned and ISO / IEC20000-1 conformant supplier, and deploys, manages and continually improves Service Management processes that are underpinned by standard technologies.

The Service Management process that Fujitsu will deploy for managing the Service will include the following key processes and functions:

- Incident Management
- Problem Management
- Event Management
- Change Management
- Availability Management
- Capacity Management
- Access to design and deployment consultants to support project.

Security and Information Assurance

Fujitsu will support the Customer accreditation process and consultancy for the SD-WAN service. This is usually included in the set-up charges however will be subject to confirmation upon agreement over Customer requirements (as detailed in the Statement of Work). Access to the Orchestration high-level design documentation will be provided to the Customer to support accreditation activity. Please note associated ITHC tests requested by the Customer to support the accreditation process will be charged using the rates from the agreed published SFIA Rate Card.

Business Continuity

Persistent data can be stored within the Service and backup is provided. No specific disaster provision is provided as part of the Service but may be specified as an optional service. Note subject to deployment option selected

Training and Consultancy

Additional consultancy can be requested from Fujitsu to support any Customer specific service management requirements, and this will be charged using the rates from the agreed published SFIA Rate Card.

4. Service Levels, for Service Packs 1 & 2

The following service levels will apply to the solution and its delivery:

- Service instance delivery target is from 16 weeks from order (using Orchestrator in Public Cloud) and advised by Fujitsu upon agreement over Statement of Works and order please note edge device delivery dates will be advised prior to order.
- Service instance delivery target of Orchestrator in Fujitsu private infrastructure / Customer infrastructure will be advised by Fujitsu upon agreement over Statement of Works.
- The Orchestration platform will be available from 99 to 99.99% availability (depending on the location of the hosted solution and the operating software deployed) during service hours for the duration of the contract. Availability covers Fujitsu service and interfaces but excludes all third-party networks and other dependent services.
- The edge device 99% to 99.99% availability depending on the device selected and options required by the Customer.

- An Appliance deployed at site availability 99% to 99.99% depending on the configuration and options required by the Customer and the operating software deployed.
- The Standard Service hours are Monday to Friday 9:00 to 17:00 (excluding public holidays). With options up to 24x7 including access to service desks.
- Premium Service is 24x7 365 days.
- Additional and bespoke support services are available on request to meet Customer requirements, including weekends and evenings.

Fujitsu may conduct maintenance activities to ensure the smooth operation of the Service and such maintenance activities may temporarily interrupt the provision of the Service. Fujitsu will use reasonable endeavours to provide the Customer with at least 14 days advance notice prior to conducting any maintenance activities, which will be scheduled outside of core service hours.

Fujitsu may temporarily interrupt the provision of the Service to conduct urgent maintenance activities. In such cases, Fujitsu will inform the Customer that urgent maintenance activities need to be undertaken and seek to agree a suitable maintenance window where possible. Such urgent maintenance activities will be excluded from availability calculations.

5. Service Delivery, for LEC Service Pack 3

Since July 2002 Fujitsu has operated for its services externally accredited European wide ISO27001 Information Security Management System (ISMS) for its datacentres, networks operations and SOC services. This includes also comprehensive processes and polices based on CPNI best practices and HMG Security Policy Framework guidance, which are incorporated with the ISO27001 ISMS, for all Customer services we deliver.

Fujitsu's ISO/IEC 27001, operations are certified by Bureau Veritas (Reference IND17.0595/UUK002399). As an example, Fujitsu implements the following best service and security practice:

- Established Data Protection Act program across Fujitsu's UK business (registration number Z6251262)
- Certified for Cyber Essentials (Reference 7288763829626211) for Fujitsu EMEA
- Certified for Cyber Essentials Plus DNS UK operations (Reference 9577641823307310)
- Fully compliant with ISO/IEC20000 IT service management
- Business continuity planning certified to ISO/IEC22301:2012
- Fully complies with ISO27002-2013 Information Technology – Security Techniques
- Operates ISO27036 – for supply chain security
- A relationship model aligned to ISO440001
- Membership and active participant in:
  - The Information Security Forum, Standard of Good Practice
  - Information Security Forum Securing the supply chain – implementation guide
  - Information Security Forum Securing the supply chain – preventing your suppliers' vulnerabilities becoming your own.
  - Active participant in UK HMG's joint best practice security working groups.

To enable certification continuity across all procedures and operations, Fujitsu carries out the Deming "Plan-Do-Check-Act" (PDCA) cycle.

Fujitsu is an ITIL® aligned and ISO / IEC20000-1 conformant supplier, and deploys, manages and continually improves Service Management processes that are underpinned by standard technologies.

The Service Management process that Fujitsu will deploy for managing the Service will include the following key processes and functions:

- Incident Management
- Problem Management
- Event Management
- Change Management
- Availability Management
- Capacity Management
- Protective Monitoring
- Access to design and deployment consultants to support project.

Security and Information Assurance

Fujitsu will support the Customer accreditation process and consultancy for the SD-WAN service, including alignment to NPIRMT assurance CNI and current NCSC guidance

Business Continuity

Persistent data can be stored within the Service and backup is provided. No specific disaster provision is provided as part of the Service but may be specified as an optional service. Note subject to deployment option selected

Training and Consultancy

Additional consultancy can be requested from Fujitsu to support any Customer specific service management requirements, and this will be charged using the rates from the agreed published SFIA Rate Card.

Secure Destruction of Edge Equipment

The Service will provide secure destruction of Edge equipment that has been withdrawn from service by the Customer or Fujitsu.

Protective Monitoring & SOC SIEM Support

The Service will provide Protective Monitoring aligned to the MITRE ATT&CK framework Service Cover 24x365 days using Elasticsearch, Kibana and Logstatsh stack. The Protective Monitoring provided shall be based on the defined Use Cases agreed with the Authority (contained in [redacted]). Changes to use cases or resources will be subject to change control. Fujitsu will provide feeds from the Elasticsearch, Kibana and Logstatsh stack, to the Authority SIEM using the format agreed in [redacted]

6. Service Levels, LEC Service Pack 3

Service Measure	Description	Support Hours	Target Response Time
Priority 1 (Major) Standard	Major business disruption: critical user or user group unable to operate, or an entire service experiencing significant reduction in system performance	Incident resolution 09:00-17:00 Mon-Fri GMT/BST excluding any public holidays with uplifts to 24x7 365 support cover see catalogue Note platform monitored 24x7 365 days. Faults may be reported 24x7 365 days	1 hours
Priority 1 (Major) Premium	Major business disruption: critical user or user group unable to operate, or an entire service experiencing significant reduction in system performance	24x7 365 days	1 hours
Priority 2 (Medium)	Partial service disruption to a live/production service.		2 hours
Priority 3 (Low)	Minor disruption: single user or user group experiencing problems, but with circumvention available.		8 hours
Priority 4 (Very Low)	Enquiry: single user or user group requiring assistance but with no direct impact on business. For example, a request for information.		24 hours

Table 6 Service Measure

All service levels have a target that 95% will be fixed within the SLA, subject to a minimum volume as advised in the Statement of Work

The incident period is measured based on the timings from ITSM; from incident raised to the time at which the incident is set to 'resolved'.

Service Desk

Fujitsu's Service Desk will provide a Single Point of Contact for incident recording, updates and resolution or fault diagnostics to your IT Staff. Fujitsu will diagnose and analyse faults and manage the fault through to resolution. Access to the Fujitsu UK Service Desk to log Incidents and Service Requests is available 24x7 via telephone or email.

Service Availability SLAs

Service	Availability Target	Availability Measurement
SD-WAN Core Platform	99.99%	24x7x365 over a quarterly period. Availability will be defined as the ability of the core platform to provide SD-WAN "Data Path Functionality between edge nodes
High Availability (HA) Sites	99.99%	24x7x365 over a quarterly period. Edge device availability will be defined as the ability for the edge HA solution to route traffic in accordance with the routing and policies programmed from the Cisco vManage to provide defined SD-WAN functionality
Standard Sites (Single edge)	99%	24x7x365 over a quarterly period. Edge device availability will be defined as the ability for the edge device to route traffic in accordance with the routing and policies programmed from the Cisco vManage to provide defined SD-WAN functionality
Non-UK Mainland Sites	97.5%	24x7x365 over a quarterly period. Edge device availability will be defined as the ability for the edge device to route traffic in accordance with the routing and policies programmed from the Cisco vManage to provide defined SD-WAN functionality

Table 7 Service Availability SLAs

Core Platform SLAs

Priority	Definition	Support Hours	Response Time	Resolution Time	Resolution Target (1) (SLA)	Resolution Target (2) (PI)
P1	**Severe business disruption:** Any existing core servers are unavailable or unable to be managed. Loss of all network connections or firewalls at core datacentres	24x7x365	30 minutes	4 hours	98% of all hardware faults fixed within 4 hours	100% of hardware faults rectified in 8 hours
P2	**Major business disruption**: Application services are unavailable or unable to be managed. Loss of a network or firewall connection at core datacentres Loss of core resilience but service to site operating	24x7x365	1 hour	8 hours	90% in 8 hours	100% in 16 hours

P3	Minor business disruption: Authority unable to manage resources within the user portal	Monday to Friday 0800-1700hrs (Excluding Bank Holidays)	5 hours	3 working days	90% in 3 working days	100% in 6 working days
P4	Minor disruption. Single user or user group experiencing problems with the user portal or API, but with circumvention available.	Monday to Friday 0800-1700hrs (Excluding Bank Holidays)	1 working day	5 working days	90% in 5 working days	100% in 10 working days
P5	Enquiry: Single user or user group requiring assistance but with no direct impacts on business. Example: a request for information or change request.	Monday to Friday 0800-1700hrs (Excluding Bank Holidays)	2 working days	10 working days	80% in 10 working days	100% in 20 working days

Table 8 Core Platform SLAs

Edge Device SLAs

High Availability Edge Sites

	Priority	Definition	Support Hours	Response Time	Resolution Time	Resolution Target (1)	Performance Indicator only
Hardware	**P1**	Hardware failure of All Edge devices preventing site connectivity.	24x7x365	30 minutes	8 hours	98% of all hardware faults fixed within 8 hours	100% of hardware faults fixed in 24 hours.
	P2	Hardware failure of a single Edge device not preventing site connectivity.	24x7x365	30 minutes	72 hours	98% of all hardware faults fixed within 72 hours	100% of hardware faults fixed in 96 hours.
Software	**P1**	Failure of All Edge devices preventing site connectivity.	24x7x365	30 minutes	4 hours	98% of all faults (where remote support capable) fixed within 4 hours	100% of faults where remote support capable) fixed within in 8 hours.
	P2	Failure of a single Edge device not preventing site connectivity.	24x7x365	30 minutes	8 hours	98% of all faults (where remote support capable) fixed within 8 hours	100% of faults where remote support capable) fixed within 24 hours.

Table 9 Edge Device SLAs – High Availability Edge Sites

Standard Edge Sites

	Priorit	Definition	Support	Response	Resolution Time	Resolution Target (1)	Performance Indicator only
Hardware	P1	Hardware failure le Edge device ng site ivity.	24x7x365	30	12 hours	98% of all hardware ed within 12 hours	100% of hardware ed in 24 hours.
Software	P1	Failure of a single vice preventing nectivity.	24x7x365	30	4 hours	98% of all faults emote support) fixed within 4 hours	100% of faults emote support) fixed within in 8

Table 10 Edge Device SLAs – Standard Edge Sites

Non-UK Mainland Edge Sites

Note: All Non-UK Mainland sites deploy a Standard edge device configuration.

	Priority	Definition	Support Hours	Response Time	Resolution Time	Resolution Target (1)	Performance Indicator only
Hardware	P1	Hardware failure of a single Edge device preventing site connectivity.	24x7x365	30 minutes	12 hours	98% of all hardware faults fixed within 12 hours	100% of hardware faults fixed in 24 hours.
Software	P1	Failure of a single Edge device preventing site connectivity.	24x7x365	30 minutes	4 hours	98% of all faults (where remote support capable) fixed within 4 hours	100% of faults where remote support capable) fixed within in 8 hours.

Table 11 Edge Device SLAs – Non-UK Mainland Edge Sites

Edge Device Replacement SLAs

Service	Support Hours	Resolution Time	Resolution Target (1)
High Availability Sites	24x7x365	72 hours	100% of all hardware replacements delivered to site within 72 hours
Standard Sites	24x7x365	72 hours	100% of all hardware replacements delivered to site within 72 hours
Non-UK Mainland Sites	24x7x365	72 hours	100% of all hardware replacements delivered to site within 72 hours

Table 12 Edge Device SLAs – Replacement SLA

Problem Management Priority Levels and Definitions

Ref	Definition	Resolution Target (1)	Performance Indicator
PM1	**Level 1 (P1):** The Problem poses significant risk to the Customers business or operations in that the Incident or series of Incidents may result in significant adverse impact to the Customer operations. No Workarounds have been identified to Resolve the Problem.	10 Working Days	20 Working Days

PM2	**Level 2 (P2):** The Problem poses no immediate risk to the Customer business or operations but may, if not Resolved, result in degradation in the performance of a Service. Workarounds are available to Resolve the Problem.	1 Month	40 Working Days
PM3	**Level 3 (P3):** The Problem poses no risk to the Customer business or operations but may in the long term impact on the overall performance of a Service	6 Months	*N/A
PM4	The percentage of all Problems occurring during the Service Measurement Period that the DJSC has Resolved within the relevant Resolution Time for each Problem. Problems shall only be considered as validly Resolved if the DJSC has Resolved the Problem and notified the Customer that it has Resolved such Problem by completing the relevant sections of the relevant Problem Record.	95%	100%

Table 13: Problem Management SLAs

Service Reporting

Service Reporting encompasses the production and delivery of defined reports to accurately report against agreed Service Level Agreements. Excluding the Daily Service Update Report, which only provides volumetric data.

Fujitsu will provide the following reports on a Daily, Monthly and Quarterly basis:

Daily Reporting

Daily Service Update Report

DJSC will:

- Generate the Daily Service Update Report each Working Day with respect to the immediately preceding Working Day (Monday – Friday)
  - Information pertaining to Saturdays and Sundays to be included within Monday's report
- Email the Daily Service Update Report to the SDM by 09:00 each day

Monthly Reporting

ITSM Monthly Service Reports (from SCSM)

- ITSM Service Reports are provided as part of the monthly service reviews and will generate the following pre-built reports on a monthly basis generated from the ITSM toolset:

Report Area	Report Name	Description
Change Management	List of Change Requests	Provides a list of change requests within a certain time frame. The data in this report includes the current status, category, and user to whom the request is assigned.
Incident Management	Incident Resolution	Provides the number of incidents, including the number of incidents past their targeted resolution time and the average time to resolution. You can filter the data by day, week, month, quarter or year.

Incident Management	List of Incidents	Provides a list of all incidents within a certain time frame. The data in this report includes the users to whom incidents are assigned, when the incidents were created, and the current status of the incidents.
Problem Management	List of Problems	Provides a list of all problems within a certain time frame.
Security Incident Management	List of Incidents	Provides a list of all security incidents within a certain time frame. The data in this report includes the users to whom incidents are assigned, when the incidents were created, and the current status of the incidents.
Configuration Items(CIs)	List of CIs	Provide a list of all CIs within the estate

Table 14 ITSM Monthly Service Reports

Additional Monthly Reports

- Forward Schedule of Change (derived via SCSM)
- Forward Schedule of Release (derived via SCSM)
- Patch Compliance Report / Statement of Conformance
- SPLA License (Service Provider License Agreement) Report
- BRA07 Spares Pool Stock Report

Quarterly Reporting

- Availability Reports

Annual Reporting

Annual AccSec Audit/Reporting; Fujitsu provides the Customer with access to records which tracks location and status of all security encrypted ACCSEC edge devices throughout its life and into its terminal state.

7. Price Approach Service Packs 1, 2 and 3

Prior to award of a contract, specialist SD-WAN solution architects will be assigned to define the functionality and to agree with the Customer the Information Assurance requirements (including hosting options). This established process will ensure the full scope of the service is understood by both parties which will be summarised in a Statement of Work. The charges applicable for the Customer will be a combination of any catalogue items (see attached) and the use of the SFIA rate card for professional services (deployment design and service elements).

License Capacity

The Customer shall be responsible for the selection of appropriate licences (for example tier capacity), which when activated shall be fixed for the Term of the agreement.

In the event the license capacity exceeds 90% for 2 consecutive months, Fujitsu will, upon 30 days notice, limit the capability to the stated license tier (thus preventing any burst capability). Upon such notice the Customer may request Fujitsu to upgrade to the next license tier for which the Customer will be responsible for the relevant license cost (including any SFIA charges to deploy). Or accept the throughput will be limited to the maximum capacity of the license tier ordered.

The Customer should note changes to a licence tier may require an Edge Device upgrade (throughput capacity), which will be advised by Fujitsu. Changes to Edge Device will be subject to charges as detailed.

Please refer to the Pricing Document.

8. Commercial Service Packs 1, 2 and 3

Ordering and Invoicing Process

The Customer will be invoiced for the Charges on a Monthly basis in arrears (based on site usage consumption model).

When remitting payment, the Customer will include the applicable Fujitsu invoice that the payment applies to.

Trial Service

A limited Trial service may be offered (subject to dialogue). Please contact Fujitsu for a demonstration of the Service.

Minimum and Maximum Terms

Unless agreed otherwise, the minimum term for the Service is twelve (24) months.

Termination terms

Thereafter Fujitsu's SD-WAN service can be terminated provided at least three (3) months' notice is given.

Consumer responsibilities

- The Customer shall, in addition to its obligations contained elsewhere in this Agreement:
- Provide and maintain list of authorised users that are allowed to contact the Service Desk
- Provide information to support the configuration of the Service, in accordance with the implementation plan developed by Fujitsu
- Provide suitably sized network connectivity and bandwidth for the delivery of Service
- Ensure appropriate network integration
- Ensure appropriate security and associated policies are in place and adhered to where required for the fulfilment of this Service
- Provide details of authorisation process and management controls and contacts for specific change and request types
- Provide appropriate contacts for escalations that are available during normal business hours.

Service constraints

- Platform Monitoring 24 hours by 365 days per year
- Incident and problem management for severity 1 & 2 incidents 24 hours by 365 days note contracted hours
- Administration of Break-fix management 24 hours by 365 days note contracted hours (options available)

- Availability event monitoring, for central infrastructure physical/environment and performance thresholds
- Capacity Management of solution
- Analytics to understand the impact of new services and users
- Service reporting including high level KPIs
- Critical patch management
- Regular firmware upgrades (to agreed policy) e.g. N-1
- Defined Service Request and Catalogue workflow Change Management.

Service exclusions

- Service transition (this is a project cost)
- Break-fix costs (engineer/part to site)
- Service catalogue pricing.

About Fujitsu

Fujitsu has been working with the public sector for over 40 years and is a global IT company offering a complete range of products, services and solutions. With expertise in digital business transformation and a proven track record of enabling Customers to digitalise with confidence, we have helped many of our Customers to digitally transform to meet changing consumer needs and take advantage of the benefits brought about by digital disruption.

Contact: government.frameworks@uk.fujitsu.com