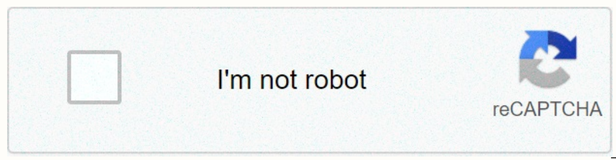

Continue

Juniper firewall design guide

Help us improve your experience. Let us know what you think. Do you have time for a two-minute survey? This section describes the computing resources, network infrastructure, and storage components needed to implement metaFabric 1.0 solutions. It also discusses the application of software, high availability, service class, security, and network management components of this solution. The purpose of the data center is to host business-critical applications for enterprises. Each role in the data center is designed and configured to ensure a high-quality user experience as possible. All functional roles in data centers exist to support applications in data centers. In the computing area, you need to choose a physical and virtual component that will host critical applications of your business, network management, and security services. These include the selection of VMs, servers, careful hypervisor switches, and knife switches. Virtual machines (VMs) are virtual computers consisting of host operating systems and applications. Hypervisor is a software that runs on a physical server, emulating physical hardware for VMs. VM operates on hypervisor emulated hardware. VM believes that it runs on dedicated physical hardware. This abstract layer allows the benefits of presentation to the operating system; regardless of changes to hardware, the operating system sees the same set of logical hardware. This allows operators to make changes to the physical environment without causing issues on servers hosted in a virtual environment, as seen in Figure 1.Figure 1: Virtual Machine DesignVirtualization also allows for unlikely flexibility on physical servers. The operating system can be moved from a set of other physical hardware with very little effort. A complete environment, to enter installed operating systems and applications, can be cloned in a virtual environment, allows for a complete backup of the environment or, in some cases, you can clone or recreate the same server on different physical hardware for redundancy or mobility purposes. This clone can be activated upon primary VM failure and allows for simple redundancy levels to exist in the data center application layer. The extension to the cloning benefits is that a new operating system can be created from this clone quickly, allowing faster service launches and faster time for revenue for new services. Servers in virtual IT data centers are just a source of physical computing that hosts VMs. Servers offers processing power, storage, memory, and I/O services to VM. Hypervisor is installed directly on the server without any type of host operating system, becoming a vulnerable metal operating system that provides a framework for data centers. Because the server hosts an income portion that generates data centers (VM and resident applications), redundancy redundancy on this layer. Virtual IT data center servers must support full hardware redundantry, management redundancy, the ability to upgrade the software during the server in service, hot swapping of power supply, cooling, and other components, and the ability to incorporate multiple servers or blade-sized servers into one logical management aircraft. The server test must be able to provide transportation between physical hardware and virtual components, connecting to the host via the 10-Gigabit Ethernet port, using 10-Gigabit Ethernet or Ethernet 40-Gigabit interface to access PODs, consolidates storage, data, and management functions, provides service classes, reduces the need for physical cables, and provides Figure 2: Server Designers seen in Figure 2, this solution includes a 40-Gigabit Ethernet connection between the excessive server node group of QFabric systems and the master's IBM Flex server house up to 14 blade servers. Other supported connection types include an oversupply 10-Gigabit Ethernet port and a 10-Gigabit Ethernet pass port. The solution also has two built-in switches per Flex server and uses MC-LAG to keep traffic flowing through data centers. The hypervisor switch is the first hop from an app server in MetaFabric architecture 1.0. The virtual machine connects to a distributed virtual switch (dvSwitch) which is responsible for mapping a set of physical network cards (pNICs) across a set of physical hosts into a logical switch that can be centrally managed by virtualization orchestration tools such as VMware vCenter (Figure 3). DvSwitch enables intra-VM traffic on the same conversion domain to pass between local VMs without leaving the blade server or virtual environment. DvSwitch also acts like a Virtual Chassis, connects various ESXi hosts at the same time, and offers port group functions (similar to VLAN) to provide access between VMs.Figure 3: Virtual SwitchThis Distributed VMware poses an attractive security challenge on the hypervisor switch, as a traditional appliance-based firewall has no vision In cases where restrictions must be placed on VM-to-VM traffic, security software can be installed on the hypervisor to perform firewall functions between VMs.Hypervisor switches are a critical piece of MetaFabric architecture 1.0. Therefore, it should support functions that enable service classes and SLA achievements. Support for IEEE 802.1p is required to support service classes. Support for parallel link aggregation (IEEE 802.3ad) is also required to ensure an excessive VM connection. As in other conversion roles, support for SLA achievements is also a necessity in this layer. The hypervisor switch should support SNMPv3, flow accounting and statistics, reflect remote ports, and centralised management reporting to ensure that the SLA can be measured and verified. To complete the configuration for a hypervisor switch, hypervisor, service classes on flow for IP storage, vMotion, management, fault tolerance, and VM traffic. As shown in Figure 4, this solution implements the following provisions for network input/output control stocks (I/O) : IP storage (33.3 per cent), vMotion (33.3 per cent), management (8.3 per cent), fault tolerance (8.3 per cent), and VM traffic (16.6 per cent). This category has been maximized for server-level traffic. Figure 4: VMware Network I/O Control DesignThe virtual IT data center has virtual equipment that is often hosted on blade servers, or servers that support multiple connectable processing blades that give blade servers the ability to host a large number of VMs. The blade server includes a power and cooling module as well as an input/output (I/O) module that allows Ethernet connection to the blade (Fig. 5). Blade conversion is done between the physical Ethernet port on the I/O module and the internal Ethernet port on the blade. In some knife servers, the 1:1 subscription model (one physical port connects to one blade) is used (this is called a conversion by route), with one external Ethernet port connecting directly to a particular blade through the internal Ethernet port. Pass models offer the benefits of allowing full bandwidth to each knife server without oversubstinguing. The disadvantage of this approach often lacks flexibility in VM mobility and provision because the VLAN interface needs to be transferred to a physical switch and a knife switch when a step is needed. Figure 5: Examples of Knife Switch, ViewAnother Mode back of blade switch operations is where the blade switch allows for an oversubscription to the knife server. In this type of blade server, there may be only 4 external ports that connect internally to 12 separate knife servers. This will result in 3:1 overs subscribed (three internal ports to each external port). The benefit to this operating mode is that it minimizes the number of connected interfaces and access changes the cable of each blade server, although oversupply link performance and connected VMs can lower the result. Although this architecture is designed for data centers that use blade servers, the design works similarly also in data centers that don't use blade servers to host VMs.Table 1 shows that both a passing blade server and an oversupply blade server are acceptable options for this solution in your data center network. In some cases, you may need faster speeds provided by the 40-Gigabit Ethernet connection to support newer equipment, while others you prefer the performance of the line rate offered by a passing switch. As a result, all three types of knife servers are supported in this design. Table 1: Blade Pass-Through and Oversubscribed Blade Server comparisons SW10G Casis SW40G Casis SWTransportYes10-Gigabit Ethernet host interface YesYesYes40-Gigabit Ethernet uplinks face to face storage, data, and management of YesYesYesClass servicesYesCable reductionNoYes (12:14)Yes (2:14)Oversubscription1:11.2:13.5:1Active/Active YesYesYesTo provide support for computing and virtualization in IT data center, virtual This solution uses:Virtual machine—VM running Windows and applications, such as Microsoft SharePoint, Microsoft Exchange, and WikiMediaServers—IBM x3750 and IBM FlexConfigure Systems cups IBM Flex servers with various ESXi hosts that support all VMs running business critical applications (SharePoint, Exchange, and MediaWiki). VSwitch configuration is distributed between various physical ESXi hosts configured on IBM servers. Hypervisor—VMware vSphere 5.1 and vCenterBlade switch—IBM EN4091 and CN4093This design for computing and virtualization segments of data centers meet the needs of this solution for workload mobility and migration for VM, location of independence for VM, VM visibility, high availability, fault tolerance, and centralized virtual switch management. This network is often the main focus of data centers because it is built to pass traffic to, from, and between application servers hosted in data centers. Given the critical role of this architecture, and various levels in the data center conversion block, it is further broken down into access conversions, aggregation conversions, core conversions, edge routing, and WAN connection. Each segment in the data center conversion role has unique design considerations related to business critical, SLA requirements, redundancy, and performance. It is in the data center changing the role of architecture that the network must be carefully designed to ensure that the purchase of your data center equipment maximizes the scale and performance of the network while minimizing costs. The access layer consists of a physical switch connected to the server and the final host. Access conversion usually focuses on implementing Layer 2 switches, but can include 3 Layer components (such as IRB) to support more robust VM mobility. Access conversion should also support high availability. In the environment of a multi-accessive or virtual wissed, where multiple physical switches can be combined to form a logical switch, redundancy can be achieved in the access layer. This type of switch architecture is built with redundancy control airplanes, MC-LAG, and the ability to upgrade individual switches while they are in service. In addition, the role of access conversion should support storage traffic, or the ability to pass data traffic via Ethernet via iSCSI and Fiber Channel via Ethernet (FCoE). Data Center (DCB) brides must also be supported by access conversion roles to enable full support of storage traffic. In DCB, support for priority-based flow control (PFC), enhanced delivery selection (ETS), and bridge exchange data (DCBX) should also be supported because these features allow properly passed storage traffic between all servers and in the data center segment. The aggregation switch acts as a thickening point between access and the core of the data center. The role of aggregate architecture works to combine a large number of smaller interfaces from access to high bandwidth stem ports that can be consumed more easily by the core switch. Redundancy should be a priority in the design of the role of aggregation as all layers of 2 flow between the data center and the core switch are combined and submitted by the role of the data center aggregation switch. At this layer, switching architecture that supports the combination of multiple switches into a logical system with the control and delivery of redundancy aircraft is recommended. This switched architecture allows redundancy features such as MC-LAG, loop-free excess routes, and software upgrades in services to enable data center administrators to meet consistently and exceed SLAs.One recommendations are to combine access and layers of aggregation of your network by using the QFabric system. Not only does the QFabric system offer a single point of allocation, management, and troubleshooting for network operators, it also collapses changing the stage for any connection to anywhere, providing lower latency, and allows all access devices only one hop away from each other, as shown in Figure 6.Figure 6: Juniper Networks QFabric Systems Enables Flat Data Center implementation , this solution uses QFX3000-M QFabric system. There are two QFabric (POD1 and POD2) systems in this solution to deliver performance and scale. QFabric POD supports 768 ports per POD and has port latency to low ports, one management point per POD, and an unseen Ethernet to support storage traffic. The use of a predetermined POD configuration allows enterprises to plan data center launches more effectively by offering predicted growth and scale in solution architecture. The main configuration steps include:QFX3000-M QFabric system configuration with 3 groups of excessive server nodes (RSNGs) connected to 2 servers of the IBM Flex System blade to deliver application traffic. The first IBM Flex System server uses the 40-Gigabit Ethernet (CNA) conversion network adapter connected to the QFabric RSNG system containing the QFX3600 (RSNG4) Node device. The second IBM Flex System Server has 10-Gigabit Ethernet through modules connected to RSNG2 and RSNG3 on the second QFabric system. Connect the EMC VNX storage platform to the QFabric system for storage access using iSCSI and NFS. Connect the QFabric system with the EX9214 core switch through a network Node group containing 2 Node devices that use four 24-port LAG configured as a dick port. OSPF Configuration POD (in the QFabric system network node pool) towards the EX9214 terrace and place this connection in Region10 as a truly degil region. Suis terrace is often configured as a Layer 3 3 device routing between multiple Layer 2 domains in the data center. The robust implementation of the core switch in virtual IT data centers will support both Layer 2 and Layer 3 to enable various interoperability and service provision in a multiple environment. Just like in the edge role, the core conversion redundancy is critical because it is also a traffic jam point between customers and applications. Properly designed data centers include a fully excessive layer of core switch that supports various interfaces (1-Gigabit, 10-Gigabit, 40-Gigabit, and 100-Gigabit Ethernet) with high density. The density of the port in the core conversion role is a critical factor because the core of the data center should be designed to support future expansion without the need for new hardware (outside the line card and interface adapter). The role of the core switch should also support a wide collection of SLA statistics, and should be aware of the services to support the service chain's statistics collection. The general location of the core conversion function in this solution is shown in Figure 7.Figure 7: SwitchingTable 2 core shows several reasons for choosing the EX9200 switch over the EX8200 switch to provide core conversion capabilities in this solution. The EX9200 switch provides a large number of Ethernet 10-Gigabit ports, support for the 40-Gigabit Ethernet port, the ability to host more analyst sessions, firewall filters, and BFD connections, and critical support for software upgrades in services (ISSU) and MC-LAG. These reasons make EX9200 change the superior option in this solution. Table 2: Core Switch Hardware - Comparison EX9200 and EX8200 SwitchesSolution RequirementEX8200EX9200DeltaLine-rate 10G128240+88%40GNoYesAnalyzer Session764+815%ACLs54K256K+375%BFD175900+415%ISSUNo (NSSU)YesMC-LAGNoYesTable 3 shows several reasons to choose MC-LAG as a delivery technology on the Virtual Chassis in this solution. MC-LAG provides dual control aircraft, uninterruptive execution, support for LACP, state replications across peers, and support for ISSU without the need for duplication engines.Table 3: Switch Forwarding Core - Comparison of MC-LAG and ChasisAttributeVirtual ChassisMC-LAGControl Planes12Centralized ManagementYesNoMaximum Chassis22ImplementationDisruptiveNon-disruptiveRequire IEEE 802.3 LACP)Replication of NoYesStateKernelCTCPRequire Dwi Routing EnginesYesNoISSUNoYesTo implements the core conversion section of the virtual IT data center, this solution uses two EX9214 switches with the following capabilities and configuration : Main features—240 Gbps line rate per slot for 10-Gigabit, Ega support for 40-Gigabit Ethernet port, 64 analyst sessions, scalable to 256,000 firewall filters, and support for dwexic submission detection (BFD), software upgrades in services (ISSU), and MCG-LA group Steps main (Rajah 8)Configuration of Layer 2 MC-LAG on/on ex9214 over QFabric QFabric F5 load balancing, and MX240 edge router (by way of excessive Ethernet link provided by the SRX3600 edge firewall) to provide the release of the route. Configure IRB and VRRP for all MC-LAG links for high availability. Configure THE IRB on EX9214 and QFabric PODs to end the boundaries of Layer 2/Layer 3. The static route configuration on the core switches to direct traffic from the Internet to the remaining load.o OSPF configuration to advertise the default route to a truly fringe area in the QFabric POD. Each QFabric POD has its own OSPF area. Also, configure the EX9214 core switch as an area-border router (ABR) connecting all three areas of the OSPF, and setting the spine area 0 over the ae20 aggregate link between the two cores of the SuisFigure 8: The core turned DesignEdge RoutingThe edge is the point in an aggregate network of all customers and the Internet connections into and out of the center Despite the high availability and redundancy is an important consideration throughout the data center, it is on the edge that they are the most important considerations the advantages serve as a choking point for all data center traffic and losses in this layer make the data center out of service. On the edge, full hardware redundancies should be implemented using platforms that support control aircraft and advance aircraft treasury, aggregation of links, MC-LAG, excessive uplinks, and the ability to upgrade software and platforms during data centers in service. This architectural role should support various protocols to ensure data centers can support any type of interconnectedness that may be offered. Side routers in data centers need support for IPv4 and IPv6, as well as ISO and MPLS protocols. Since data centers may be multiple tenants, the vast range of routing protocols should also be supported, to include static routing, RIP, OSPF, OSPF-TE, OSPFv3, IS-IS, and BGP. With a large-scale multi-tenant environment, it is important to support Virtual Private LAN Services (VPLS) through bridge domain support, overlapping VLAN IDs, integrated routing and bridges (IRB), and IEEE 802.1Q (QinQ). The advantages should support a complete set of MPLS VPNs, including L3VPN, L2VPN (RFC 4905 and RFC 6624, or Draft Martini and Kompella, respectively), and VPLS.Network Address Translation (NAT) are another factor to consider when designing the data center's edge. It is likely that some of the customers included by the data center will have a private network address scheme overlapping. In an environment where Internet access directly to data centers is enabled, NAT is required to translate public IP addresses that can be redired to the private IP addresses used in data centers. Advantages must support NAT Base 44, NAPT44, NAPT66, Twice NAT44, and Finally, because the edges are the entry points and egress of data centers, execution should support robust data collection to enable administrators to verify and prove strict service-level agreements (SLAs) with their customers. Their. The edge layer should support the average accumulation of traffic flow and statistics, and should at least support the ability to report the exact traffic statistics to include the exact number of bytes and packets received, sent, queued, lost, or dropped, per application. Figure 9 shows the location of the edge routing function in this solution. Figure 9: The role of Edge RoutingWANThe WAN provides transportation between end users, isolated sites of enterprise, and data centers. There are several different WAN topologists that can be used, depending on the business needs of the data center. Data centers can only connect directly to the Internet, using easy IP-based access directly to servers in data centers, or secure tunnel approaches using generic routing encapsulation (GRE) or IP Security (IPsec). Many data centers serve a wide base of customers and favor Multiprotocol Label Switching (MPLS) connections via the MPLS network managed service providers, allowing customers to connect directly to the data center through the carrier's MPLS spine. Another approach to WAN is to allow direct peers between customers and data centers; This approach allows customers to bypass transit peer links by establishing a direct connection (for example, via a leased personal line) to the data center. Depending on the business needs and performance requirements of the app-hosted data center, WAN's connection options offer the first option in determining the performance and security of data center applications. Choosing a private peer or MPLS connection offers better security and performance at higher expenses. In cases where the hosted application is not sensitive to security and performance, or where the application protocol offers built-in security, a convenient Internet-connected data center can offer a suitable level of security and performance at a lower cost. To implement the edge routing and WAN parts of the virtual IT data center, this solution uses the MX240 Universal Edge router. Since the MX240 router offers both a Routing Engine and isSU at an affordable price point, it is the preferred option over the smaller MX80 router. The main connection and configuration steps are: Connect the MX240 suburb router to the service provider's network to provide Internet access to the data center. Configure two edge routers to become an EBGP peer with 2 service providers to provide excessive Internet connection. Configure IBGP between 2-edge routers and uses its next self-export policy. Configure BGP local preferences on key service providers to offer optional exit points to the Internet.Export dynamic, condition-based routes, default to the Internet to OSPF on both edge routers to the Internet edge firewall and terrace to provide Internet access for virtual IT data center devices (Rajah 10). Configure both sideblocks in Region 1 for OSPF. Power Network Address (NAT) to change the private IP address to a public IP address. Figure 10: Edge Routing DesignThis design for data center network segment meets the requirements of this solution for 1-Gigabit, 10-Gigabit, and the 40-Gigabit Ethernet port, constant data and storage, load balancing, experience quality, network segments, isolation and separation of traffic, and time synchronization, and synchronization time The role of metaFabric architectural storage 1.0 is to provide files and block the storage of data so that all hosts can Data storage can be local to VM, such as databases that live in hosted, or shared applications, such as MySQL databases that can be on various storage to serve a variety of different applications. MetaFabric 1.0 architecture requires the use of shared storage to enable computing virtualization and VM mobility. One of the main goals of virtual IT data centers is to devote both data and storage to the same network infrastructure to reduce overall costs and make operations and troubleshooting easier. There are several different options when gathering storage traffic: FCoE, NFS, and iSCSI. One of the latest trends in building green field data centers is to use IP storage and deliberately choose not to integrate the legacy Channel Fibre network. In addition, since iSCSI has better performance, lower write read response times, lower costs, and full application support, iSCSI offers better storage network options on NFS. In addition, storage traffic is very short and falling sensitive, so it is very important that network infrastructure provides non-lost Ethernet services to prioritize all storage traffic correctly. As a result, this solution uses both iSCSI and NAS for storage, and provides unseen Ethernet services to guarantee delivery of storage traffic. Table 4 shows comparisons of FCoE, NFS, and iSCSI. Because NFS and iSCSI meet the same requirements provided by FCoE, plus the ability to scale to 10-Gigabit Ethernet and beyond, NFS and iSCSI storage protocols are the top choices for MetaFabric 1.0 solutions. Table 4: Comparison of Storage Protocol NeedsCOENFSiSICSILossless EthernetYesYes10GE and beyondNoYesYesConverged data and storageYesYesYesYesLess From end-to-end Latnation 3pYesYesFigure 11 shows the storage traffic path as it moves through data centers and highlights the benefits of queuing priorities to provide unseen Ethernet transport for storage traffic. By configuring Preferences Flow Control (PFC), storage devices can monitor storage traffic in VLAN Storage and notify servers when traffic jams occur. can pause send additional storage traffic until after the storage device has cleaned the crowded recipient's thickener. However, other lines were unaffected and unexpected traffic continued to flow uninspired. Figure 11: Missing Ethernet Design Packet Flow Storage for storage is as follows:The server sends storage traffic to the QFabric system. The QFabric system classifies traffic. Traffic is enforced in priority. The QFabric system sends traffic to various storage. Various storage receives traffic. Various storage sends traffic back to the QFabric system. The QFabric system classifies traffic. Traffic is enforced in priority. The QFabric system sends traffic to the server and VMs.Server receives traffic. To implement the virtual IT data center storage section, this solution uses EMC VNX5500 unified storage with multiple single storage. This storage is connected to the QFabric POD, which in turn connects to servers and VMs, as seen in Figure 12. The design assumes that data center architects want to save costs initially by sharing a variety of single storage with various QFabric PODs. However, the design can evolve to allocate a storage array for each QFabric POD, as usage and demand guarantee such expansion. Figure 12: The Storage DesignThis Solution also implements the Data Center Bride (DCB) to enable full support of storage traffic. In DCB, support for priority-based flow control (PFC), enhanced transmission selection (ETS), and The Exchange Bridging Capability Data Center (DCBX) allows storage traffic to pass correctly between all servers and storage devices in the data center segment and to deliver an unseen Ethernet environment. This design for virtual IT data center storage segments meets the requirements of this solution for scale, unseen Ethernet, the ability to boot from shared storage, and support for multiple protocol storage. Applications in virtual IT data centers are built as Virtual Machines (VMs) and hosted on servers, or physical computing resources that are on the blade server. The design for this app meets the requirements of this solution for business critical applications and high performance. MetaFabric Solutions 1.0 supports a complete heap of software that includes four main application categories: compute management, network management, network services, and business critical applications (Figure 13). The app runs on IBM servers and VMware vSphere 5.1.Figure 13: Virtualized IT Data Center Solution Software StackVMware vCenter is a virtualization management platform that offers centralized control and visibility into computers, storage, and network resources. Data center operators use de facto, vCenter industry standards daily to manage and allocate VMs. VMware vCloud Director allow data center managers to create internal cloud services and division of virtualization environment into segments that can be administered by a separate business unit or administrative entity. Resource groups can now be divided into data centers that can offer their own independent virtualization services. The use of vCenter and vCloud directors offers support element of the software application for MetaFabric 1.0 solution. MetaFabric Solutions 1.0 uses the Junos Space Management App to provide network allocation, orchestration, and inventory management. The app includes Network Director for wired and wireless data center network management, and the Security Director for security policy administration. Network load balancing is a common network service. There are two methods for providing network load balancing: virtual-based and hardware. Virtual load balancing operates in hypervisor as a VM. One of the benefits of balancing virtual load is the rapid provision of extra load balancing power. Another benefit is that the administration of virtual load balance can be delegraded to other administrative entities without affecting other applications and traffic. However the downside to virtual load balancing is that performance is limited to the number of computing resources available. The remaining hardware loads offer more performance in traffic signage and encryption and decryption of SSL with dedicated security hardware. MetaFabric Solution 1.0 using local traffic managers (LTM) from F5 Networks.Load balancing provides the following services:Advertising the existence of trafficDistrbute applications across a set of servers. Features of leverage such as acceleration and compression of SSL. Set up additional Layer 7 features. Software applications are made of various server levels; the most common are Web Servers, applications and databases. Each server has its own set of scripted responsibilities. Web Level handles interactions with users and applications. The application level handles all application logic and programming. The database stage handles all data storage and application inventory. The following software applications have been tested as part of a MetaFabric 1.0 solution:SharePointThe SharePoint's Microsoft application requires three levels: Web, application, and database. The Web Level uses Microsoft IIS to handle Web tracking and interaction with end users. The application stage uses Microsoft SharePoint and Active Directory to provide file sharing and content management software. Finally, the database stage uses Microsoft SQL Server to store and organize application data. The Microsoft ExchangeThe Exchange app requires two levels: the Web level, and the second stage that combines applications and databases into one stage. MediaWikiThe MediaWiki app requires two levels: a Web affiliate and an app level, and a database stage. Apache httpd is combined with hypertext preprocessors (PHP) to create and present applications, while data is stored at the database level with MySQL.This design meets the availability requirements redundancy hardware and relevance software redundancy. To set up hardware redundancy in virtual IT data centers, this solution uses:Excessive server hardware—Two overs—Two 3750 standalone servers and two IBM Pure Flex System ChassisRedundant access and aggregation PODs—Two QFX3000-M QFabric systemsRedundant core switches—Two EX9214 switchesRedundant edge firewalls—Two SRX3600 Services GatewaysRedundant edge routers—Two MX240 Universal Edge routersRedundant storage—Two EMC VNX5500 unified storageRedundant load balancers—Two F5 LTM 4200v load balancersOut-of-band management switches use Virtual Chassis technology—Four EX4300 switchesTo provide software redundancy in the virtualized IT data center, this solution uses:Graceful restart—Helper routers assist restarting devices in restoring routing protocols, state, and convergence. Graceful Reblocking Engine Redemption—Make sure the state of the operating system is moved between the parent and the Engine Blocker backrest in the Juniper Circuit device. Upgrade software in service (for terrace and edge blockers)—Allowing network operating systems to be upgraded without downtime. MC-LAG—Allows aggregated Ethernet interface packages to contain faces from more than one device. Non-stop logging—Make sure that the Condition of Layer 3 protocol is moved between the parent and the Engine Blocker backrest.Nonstop bridge—Make sure that the Operation Layer 2 protocol state between the parent and the Engines.Nonstop Docking backrest increases the software level—(for QFX3000-M QFabric system POD)—Allows the network operating system to be upgraded to the minimum effect to be stated. CyberBlock redundancy protocol (VRRP)—Provides a virtual IP address for traffic and advances traffic to one of the two peer routers, depending on which one operates. MC-LAG Design ConsiderationsTo allows all links to traffic without using the Primary Spanning Protocol (STP), you can configure MC-LAG on the edge eraser and terrace suis. The edge blocker uses MC-LAG towards the edge firewall, and the terrace suis uses MC-LAG in the direction of each QFabric POD, application load balancer (F5), and outer management band (OOB). The multichassis linking (MC-LAG) pool is a feature that supports aggregate Ethernet packages spread across more than one device. The Link Aggregate Control Protocol (LACP) supports MC-LAG and is used for dynamic configuration and monitoring of the link. Options available for MC-LAG include On/Ready (where one device is active and the other helps if the active device fails) or On /On (where both devices actively take part in an MC-LAG connection). For this solution, Mc-LAG Active/Active takes precedence as it provides link level protection and nod rating for layer 2 and Layer 2/ Layer 3 hybrid environments. Active/ActiveMC-LAG Active/Active MC-LAG highlights have the following characteristics:Both terraces have an active aggregate Ethernet expert interface and advance traffic. If one of the suis terrace fails, suis the other terrace advance traffic. Traffic is loaded by default, so the efficiency of the link level is 100 percent. The Active/Active Method has been concentration from the Active/Standby method. Rapid convergence occurs because information is changed between routers during operation. After the failure, the remaining core switches of the operation do not need to release any routes and continue to advance traffic. Routing protocols (such as OSPF) can be used through the MC-LAG/IRB interface for termination of Layer 3. If you configure Layer 3 protocol in the core, you can use the routing interface and integrated bridges (IRB) to offer layer 2 and Layer 3 environments on the core switch. Active/Active also offers maximum use of resources and end-to-end load balancing. To extend the aggregate batch of links (LAG) across two devices (MC-LAG):Both devices need to synchronize their aggregated LACP Ethernet configurationLearned MAC addresses and ARP entries must be synchronized. The MC-LAG requirements above are achieved using the following protocols/mechanisms as shown in Figure 14:Interchassis Control Protocol (ICCP)Interchassis Link Protection Link (ICL-PL)Figure 14: MC-LAG – ICCP and ICL Design ICCPICCP is the control plan protocol for MC-LAG. It uses TCP as a transport protocol and Ddirectional Submission Detection Detection (BFD) for fast concentration. When you configure an ICCP, you must also configure BFD. The ICCP synchronizes the configuration and operating conditions between the two MC-LAG peers. The ICCP also synchronizes mac addresses and ARP entries learned from one MC-LAG node and share it with other peers.• Peers with ICCP partner loopback IP addresses are recommended to avoid any failure of direct links between MC-LAG peers. As long as the logical relationship between peers remains, the ICCP remains. While you can configure ICCP either a single link or an aggregate package link, the Aggregated LAG Ethernet is prioritized.• You can also configure ICCP and ICL links on an Aggregate Ethernet tie under multiple logical interfaces using flexible VLAN markings supported on the MX series. ICL-PLICL is a special 2 layer link for Active-Active only between MC-LAG peersICL-PL required to protect MC-LAG connectivity in case the failure of all cores faces a link that matches one MC-LAG node. If the traffic receiver is the sole house to one of the MC-LAG (N1) nodes, ICL is used to advance the accepted packets by means of the MC-LAG interface to another MC-LAG (N2) node. Split horizons are enabled to avoid loops on ICLThere no MAC LEARNING data aircraft over the ICL.MC-LAG Special Configuration Group ParametersRedundancy GROUP ID—ICCP uses redundancy groups to associate various cassualties performing the same redundancy function. Redundancy groups set up communication channels so that applications in ICCP colleagues can reach each other. Redundancy group ID is similar to that of an identifier Mesh. MC-AE ID—Multiple chassis Ethernet (MC-AE) IDs are between faces per casis. For example, if one mc-ae face is spread across several terraces, you should provide the same redundancy ID. When an application wants to send a message to a specific redundancy group, the app provides information and the ICCP communicates it to members of the redundancy group. Service ID—A new service ID object for the bridge domain overrides any global switch options configuration for the bridge domain. Service ID is unique across the network for services provided to enable proper synchronization. For example, service ID synchronizes applications such as IGMP, ARP, and MAC learning addresses for specific services across core switches. (Note: Both MC-LAG peers must share the same service ID for the given bridge domain.) MC-LAG Active / Active Layer 3 Routing FeaturesMC-LAG Active / Active is a Logical Link Layer 2. The IRB interface is used to create integrated Layer 2 and Layer 3 links. As a result, you have two design options when assigning IP addresses across MC-LAG peers:Option 1: Active/Active MC-LAG VRRP provides the same virtual IP address and virtual MAC and unique physical IP and MAC addresses. Both types of addresses are required if you configure the routing protocol on the MC-LAG Active/Active interface. VRRP data submission logic has been modified in Junos OS if you configure both Active/Active MC-LAG and VRRP. When configured simultaneously, both MC-LAG and VRRP colleagues forward traffic and traffic load balance between them, as shown in Figure 15.Figure 15: VRRP and MC-LAG - Active/Active OptionData Packets received by peer VRRP backups on MC-LAG member links are forwarded to the core links without sending them to VRRP Option 2: Mac synchronization address Figure 16 provides a unique IP address of each peer, but shares the MAC address You need to use option 2 if you do not intend to configure routing protocols on the MC-LAG Active/Active interface. Figure 16: MC-LAG – Mac OptionYou Synchronization Address configures the same IP address on the IRB interface of both nodes. The lowest MAC address was selected as the entry MAC address.• Peers with higher IRB MAC addresses learned peer MAC addresses via ICCP and installed peer MAC addresses as their own MAC addresses. On the MX Series platform, configure mcae-mac-sync in the bridge's domain configuration. On the EX9214 switch, the mcae-mac-sync configuration in the VLAN configuration. We recommend Option 1 as the preferred method for MetaFabric 1.0 solution for the following reasons:The solution requires OSPF as a routing protocol between QFabric POD and the core of reviving the IRB's MC-LAG interface and only Option 1 supports routing protocols. Layer 3 is extended to QFabric PODs for several VLAN for Hybrid Layer 2/Layer 3 connections to the core. MC-LAG Traffic Forwarding Rules 17: MC-LAG – Traffic Delivery Rules shown in Figure 17, the following transmission rules apply to MC-LAG Active/Active:Traffic received on N1 from MCAE1 can be flooded with ICL links to reach N2. When it reaches N2, it cannot be flooded back accepted on SH1 can be flooded with MCAE1 and ICL by N1 way. When N2 receives SH1 traffic across ICL links, it can no longer be flooded with MCAE1. N2 also received SH1 traffic via mc-AE links. When receiving packets from an ICL link, the MC-LAG peer submits traffic to all local SH links. If a matching MCAE link on peers is down, the recipient's peers also forward traffic to the MCAE link. Note: The ICCP is used to signal the state of the MCAE link between peers. When N2 receives traffic from ICL links and N1 core links, traffic cannot be forwarded to N2 core links. MC-LAG Active / Active High Availability EventsICCP down, when ICL goes up:Figure 18: MC-LAG – ICCP Down Here is an action that occurs when the ICCP link goes down and the ICL link is:By default, if the ICCP link fails, as shown in Figure 18, the default peer to the local LACP system ID and the link for only one peer (for which one consults with the router [CE] the client first) is attached to the package Until the LACP gathers with a new system ID, there will be minimal traffic effects. One peer remains active, while the other enters standby mode (but this is nondeterministic). The access switch selects a core switch and sets up a LACP peer. To optimize this situation, include an active statement of priority status control on active peers. With the dominance of priority status control configured on active peers, peers remain active and retain the same LACP system ID. With the force-icl-down statement, the ICL link was closed when the ICCP link failed. By configuring these statements, the effects of traffic are minimized during the failure of the ICCP link. ICCP up

and ICL down:Figure 19: MC-LAG – ICL DownHere is an action that occurs when ICCP links go up and ICL links go down:If you configure peers with ready-to-run control statements, the MC-AE interface shared with peers and connected to ICL is down. This configuration ensures a loop-free topology as it does not advance the wipest packets in the Layer 2 network. Active MC-LAG node goes down with ICCP loopback peer with priorities-control active on both peers:Figure 20: MC-LAG - Peer DownHere is an action that occurs when both MC-LAG peers are seized configured with an active statement of optional status control and active peers go down:When you configure MC-LAG Active/Active between SW1/SW2 and POD QFabric, SW1 becomes active and SW2 becomes standby. During the ICCP failure event, if SW1 has an active statement of optional status control and it fails, SW2 is not aware of the failure of the ICCP or SW1. As a result, SW2 mcae-id switches to the default LACP system ID, which causes the MC-LAG link to go down and up, resulting in a long traffic reconciliation time. To avoid this, configure the priority status control active statements on both SW1 and SW2. Also, you must prevent ICCP failure by configuring ICCP on between coil faces. Configuration on both active and ready-made peers. BFD helps detect peer failure and allows for sub-second reconstruction. The design for high availability in MetaFabric 1.0 solution meets the need for hardware redundantcy and software comparison. Key design elements for service classes in this solution include network control (OSPF, BGP, and BFD), virtualization control (high availability, fault tolerance), storage (iSCSI and NAS), business critical applications (Exchange, SharePoint, MediaWiki, and vMotion) and best venture traffic. As seen in Figure 21, the entry packets are arranged, tasked with queuing up based on the type of traffic, and delivered based on the importance of traffic. For example, iSCSI's non-missing Ethernet traffic has the largest line and highest priority, followed by critical traffic (high offense tolerance and availability), business critical application traffic (including vMotion), and bulk best-effort traffic with the lowest priorities. Figure 21: Service Class – Classification and Queue As seen in Figure 22, the following percentages are allocated for service classes in this settlement: network control (5 per cent), virtualization control (5 per cent), storage (60 per cent), business critical applications (25 per cent) and best venture traffic (5 per cent). This category has been maximized for network-level traffic, as the network supports multiple servers and switches. As a result, application storage and traffic traffic is the most critical type of traffic in the network, and this provision has been confirmed by our testing. Figure 22: Service Class - Buffers and Transmitter Design To provide service classes in virtual IT data centers and meet design requirements, this solution uses:Ethernet without loss for storage traffic allocation to distinguish Traffic traffic for critical applications of businessBest-traffic efforts for data trafficSecurity is an important component of any network architecture and virtual IT data center On the perimeter, security is focused on getting the edge of data centers from external threats and by providing secure entrances to the Internet. Remote access is another area where security is important in data centers. Operators often require remote access to data centers to perform new maintenance or activation of services. These remote access must be guaranteed and monitored to ensure that only authorized users access. Established verification, authorization and accounting (AAA) mechanisms must be prepared to ensure only authorised operators are allowed. Since data centers are cost and income centers that can house critical data and applications of many different enterprises, Various factors are an absolute need to get remote access correctly. The security of software applications in virtual IT data centers is security provided between VMs. A A A communication agreement between VM takes place at data centers and controlling these interactivity is an important security concern. If the server is supposed to access a database that is on another server, or on a storage setting, the virtual security appliance should be configured to restrict communication between such resources to allow only the necessary protocols for operation. Limiting inter-source communication prevents security breaches in data centers and may be a requirement depending on the regulatory requirements of hosted applications (HIPPA, for instance, can determine which security protections must exist between patient and business data). As discussed in the Virtual Machines section, security in virtual networks, or between VMs, contrasts with the security that can be implemented on physical networks. In physical networks, hardware fireins can connect to different subnets, security zones or servers and provide security between those devices (Fig. 23). In virtual networks, physical fire walls do not have the ability to see traffic between VMs. In this case, the virtual hypervisor security appliance needs to be installed to enable security between VMs.Figure 23: Physical Security Compared to Virtual Network SecurityWhen getting a VM, you need a comprehensive virtualization security solution that implements hypervisor security with full inspection; including high-sized hypervisor-based state fire walls; using an integrated interruption detection system (IDS); provide virtualization-specific antivirus protection; and offers unmatched scalability to manage the security of multiple cloud data centers. Juniper Networks Firefly Host (formerly vGW) offers all these features and allows operators to monitor software, patches, and files installed on VM from the central location. Firefly Host is designed to be centrally managed from the display of one pane, giving administrators a comprehensive view of virtual network security and VM inventory. Table 5 shows the relative merits of three application security design options: vSRX, SRX, and Firefly Host. Because of other options of lack of detection and prevention of disruption, quarantine capabilities, and performance and scale of mission critical line rates, Firefly Host is the preferred option for this solution. In addition, Firefly Host is integrated into all VMs and provides every final point with its own virtual firewall. Table 5: Application Security OptionsRequirementvSRXSRXFirefly HostStateful security policiesYesYesCentralized managementYesYesIntrusion detection and preventionYesQuarantineNoNoYes10G online rate performance at scaleNoYesTo provides application security in virtual IT data centers, This solution uses Juniper Networks Firefly Host to provide security of VM-to-VM applications. Host integrates with VMware vCenter for comprehensive VM safety and maintenance. Rajah 24: DesignIn Rajah 24's Request for Salvation, the following VM-to-VM traffic:A VM sends traffic to VM.Perkas destination Host Firefly checking traffic. Traffic corresponds to security policies. The ESXi host sent traffic. The second ESXi host received traffic. Fireflies host checks traffic. Traffic corresponds to security policies and allows traffic to go straight to the destination. VM destinations receive traffic. The edge firepower operates security functions such as Network Address Translation (NAT), detection and prevention of intrusions (IDP), security policy enforcement, and virtual private network services (VPN). As shown in Figure 25, there are four locations where you can provide security services for physical devices in your data center:Firewall filters in the QFabric PODsFirewall filter system in the switchesDedicated core, State firewalls (such as SRX3600)Physical firewalls connected to the QFabric PODsFigure 25 system: Physical Security DesignThis Solution performs option 3, which uses state firewalls to protect the flow of traffic moving between edge router and core switch. Anything below POD level is protected by the Firefly Hosting app. To provide perimeter security in virtual IT data centers, this solution uses the SRX3600 Service Gateway as a firewall edge. This firepney offers up to 55-Gbps of fire wall performance, which can easily support the VM traffic generated by this solution. The main configuration task includes:Configure the SRX gateway as an active/backup cluster. Place the excessive Ethernet reth1 group (configured towards the edge router) in the non-trust zone. Place the reth0 (configured towards the core switch) in the trust zone. Configure security policies for traffic coming from non-trust zones to allow only access to data center applications. Resource Network Translation Configuration (SNAT) for Internet access to application servers (private addresses) to provide Internet access. Destination Network Address Translation Configuration (DNAT) for remote access to data centers by translating the Pulse gateway's internal IP address to an Internet-accessible IP address. The side fireprint configuration in the OSPF area 1.MetaFabric Solution 1.0 requires secure remote access into the data center environment. Such access must provide multifactor authentication, granular security controls, and the scale of the user giving multitenant data centers the ability to provide access to administrators and access to thousands of users. Secure remote access applications must be accessed over the Internet; capable of providing encryption, RBAC, and two-factor authentication; access virtual environment; and scale to 10,000 users. Table 6 shows the MAG entrance comparison and the Junos Pulse gateway options. For this solution, junos Pulse entrance superior as it offers all MAG entrance capabilities as well as being a virtual app. Table 6: Data Center Remote Access GatewayVirtual Pulse GatewayInternet accessibleYesYesEncryptionYesTwo-factor confirmationYesYesScale to 10,000 usersYesVirtualizedNoYesTo provides secure remote access to and from virtual IT data centers, This solution uses Juniper Networks SA Series SSL VPN Appliance as a junos Pulse remote access system and gateway. Figure 26: Remote Access Flowers shown in Figure 26, remote access flows in virtual IT data center occur as follows:Users log in from the Internet.The user session is routed to the fireprint. NAT destinations were performed during the session. Authorized users match security policies. Traffic is forwarded to the entrance of Junos Pulse. Traffic arrives on an untrue interface. Trusted traffic allows local addresses to be assigned to users. Users are verified and granted access via RBAC. The design for security in the MetaFabric 1.0 solution caters to perimeter security, application security, and secure remote access. Network management is often reduced to its basic services: errors, configuration, accounting, performance, and security (FCAPS). In virtual IT data centers, network management is more of a simple tool that simplifies FCAPS: it's an enabler of growth and innovation that provides end-to-end orchestration of all data center sources. Effective network management provides the display of one data center pane. This one-pane display allows visibility and mobility and allows data center operators to monitor and change the environment across all levels of the data center. Network management in virtual IT data centers can be broken down into seven stages (Fig. 27). Figure 27: Seven Network Management Level Model is a combination of this stage that provides complete orchestration in the data center and allows operators to quickly turn on new services, and change or resolve existing service problems using the display of one data center pane. The user interface is responsible for interacting with data center operators. This is the interface from which the data center pane view is presented. From the user interface, operators can view, modify, delete, or add new network and service elements. The user interface acts as an enforcement point of a single role-based access policy (RBAC), allowing operators to access all authorized devices while protecting other resources from unapproved access. The application programming interface (API) enables the management of one pane by providing the same interface and language to applications, support tools and other devices in the data center network (REST API is an example commonly used in network management). THE API enables the display of one pane by absorbing all elements and presenting it through a single network management interface - the user interface. The network management platform should have the ability to support specific applications. Applications on the network space is specifically designed to solve specific problems in the management of the data center environment. A single application on the network management platform can be responsible for configuring and monitoring security elements in data centers, while other applications are designed to manage physical and virtual conversion components in data centers. Again, the abstract of all these applications into the one-pane view is essential for data center operations to ensure simplicity and similar management points at the data center. The next stage of data center network management is the global network display. In short, this is the stage where a complete view of the data center and its resources can be installed and viewed. This layer should support topological discovery, automatic discovery not only of the device, but how those devices intertple with each other. Global network views should also support the calculation of routes (the link distance between the network elements as well as the set of routes established between those network elements). The virtualization stage of network management resources allows multi-endpoint management in data centers and acts as an abstract layer that allows operators to manage endpoints that require different protocols such as OpenFlow or Device Management Interface (DMI). The level of general data services of network management enables various applications and interfaces on network management systems to share relevant information between layers. Applications that manage a set of endpoints may require network topological details to map and potentially reject changes to those network devices. This requires applications in data sharing network management systems; this is enabled by a normal layer of data service. Managed devices in network management roles are only endpoints managed by network management systems. The device includes a physical and virtual switch, router, VM, blade server, and security equipment, to name a few. Managed devices and orchestration services between such devices are the primary purpose of the network management system. Network management should be the answer to questions, how does the data center operator easily stand up and maintain services in data centers? The network management system dredges the execution and operation of managed devices in theFinally data center, the integration adapter is required in a complete network management system. Since every device in a data center may not be managed by a single network management system, other equipment or services may be required to manage the entire data center. The integration and synchronation of these various network management tools is the purpose of this layer. Some elements of data centers such as Virtual Machines may require an ESXi VMware server to manage VM and hypervisor switches, while networks the appliance monitors the environmental conditions and performance on the host server. The third system may be responsible for configuring and monitoring the network connection between the knife server and the entire data center. The integration adapter allows each of these components to speak to each other and, in many cases, allows a single network management system to control the entire network management footprint from one glass pane. Requirements for off-band management include:Administration of computing segments, networks, and data center storage. Isolation of control aircraft from data aircraft so that management networks remain accessible. Support for Ethernet's 1-Gigabit management interface. Provides traffic separation across computing, networking and storage segments. Enable administrator access to management networks. Deny management traffic to management. Some key elements of this design are seen in Figure 28.Figure 28: From Network Management Band DesignTo providing out-of-band management in virtual IT data centers, this solution uses two pairs of EX4300 switches configured as Virtual Casis (Figure 29). Key connection and configuration steps include:• Connect all OOB network devices to the EX4300 Virtual Casis (100-Megabit Fast Ethernet and 1-Gigabit Ethernet). Configuration of OOB Casis Virtual Management System EX4300 in OSPF 2.Connect self-server 2 IBM 3750 that hosts VM management (vCenter, Junos Space, Network Director 1.5, domain controller, and junos Pulse gateway) to Ex4300 Virtual Chassis.• Create four VLAN to separate storage, calculations, networks, and management traffic with each other. Manage and monitor VM on test beds using VMware vSphere and Network Director 1.5.Figure 29: From Band Management - DetailTo provides configuration and network allocation in virtual IT data centers, this solution uses the Juniper Network Director. Network 1.5 directors are used to manage network configurations, provisions and monitoring provide security policy configurations at virtual IT data centers, this solution uses the Juniper Network Security Directorate. The Director of Security is used to manage the configuration and provisions of security policies. This design meets the requirements of network management to manage both virtual and physical components in the data center and handle FCAPS considerations. The solution must support 20,000 virtual machines and increase up to 2,000 servers. The solution must support a total of 30,000 users.10,000 Microsoft Exchange10,000 Microsoft SharePoint users transactions10,000 mediaWiki transaction users The solution must offer less than 3µ latency between the server and the latency 21µ between PODSThe solutions must provide high availability. Less than one second convergion*No single point of FailurePublished: 2015-04-20 2015-04-20