DISASTER RECOVERY PLAN

SAMPLE RUNBOOK

SUMMIT

Customer (Customer#)

This Runbook has been created to serve as the foundation for [COMPANY]'s managed infrastructure and DRaaS services.

This document captures all critical information for these two (2) infrastructure solutions, including system requirements and business processes.

This is a living document which will be kept up to date by both parties – with all refinements pending joint approval from [COMPANY] and Summit.

Managed infrastructure / private cloud

This Runbook has been created to serve as the foundation for [COMPANY]'s managed infrastructure services. This document captures all critical information for this infrastructure solution in this section, including system requirements and business processes.

Read first

- Any changes to the DC1 stack must be reflected in DC2 and the appropriate DR procedures updated if applicable.
- Any changes to firewall or switch configurations in DC1 must be reflected in the DC2 environment and the appropriate DR procedures updated if applicable.
- DC1 refers to the DC1 environment and devices. DC2 refers to the DC2 environment and devices

Switches

- 2 x MODEL Virtual Chassis in DC1
- Name: sw1.DC1
- Cab: ----
- 1 x MODEL in DC2

Network diagram

- File: [COMPANY] Network Diagram 1.pdf

Physical layout

DC1 Hardware

- 2 Juniper MODEL configured as an active passive pair. Interconnection is per Juniper standard protocol.
- 2 Juniper MODEL configured as a virtual chassis using 40GBase interconnects per Juniper standard. Both physical chassis are configured as type Routing Engine.
- The switch and firewall complexes are cross connected using SFP DAC cables with four connections providing two logical channels so any combination of two chassis can fail without losing connectivity
- The MODEL cluster is connected to the Summit network with one 10G fiber link from each physical switch to a pair of Juniper MODEL routers, which then connect via robust means the Summit network.
- External access is provided by MODEL SSL gateway connected to the MODEL cluster for both internal and external sides.

DC2 Hardware

- One Juniper MODEL is cross connected to one Juniper MODEL using a redundant link. There are two 10G fiber uplinks from the MODEL to the Summit network, one to each of the DC2 core MODEL routers.
- External access is provided by MODEL SSL gateway connected to the MODEL for both internal and external sides.

Other Physical Considerations

- There is a leased line provided by VENDOR going from the CITY office to DC1 which is terminated in the MODEL cluster. This is a linear connection, alternate access in the case of failure or maintenance is provided via encrypted tunnels.

Logical interconnections

The logical connectivity for both sites is identical, with the links from the MODEL core to the MODEL facility and from the MODEL facility to the firewalls is done using 802.1Q trunks. Detailing the trunks and VLAN lists is out of the scope of this document.

Since the physical links to the Summit network terminate on the MODEL, IPv4 and IPv6 transit connectivity is delivered to the firewall complexes are simple layer 2 tags over the interconnecting trunks. Using discrete tags in this fashion allows us to simplify the physical maintenance and also allows us to monitor traffic levels for each function and present them in the Summit portal.

DC2 and DC1 are logically interconnected using an layer 2 virtual circuit provided by Summit. At this writing this is provisioned using MPLS/VPLS and although this could change in the future as the Summit network evolves over time Jumbo Frame support is provided on the virtual circuit.

DC1 is logically connected to several locations using pre-provisioned IPSec tunnels using shared passwords for <redacted>.

Routing

Since there are only two direct endpoints the routing requirements for this project are very straightforward: get from one end to the other. We are following best practices of segregating customer traffic from network management traffic, and since there are a significant number of separate bits of hardware to be connected, we are also following industry best practices for routing protocols.

All management traffic for the Juniper hardware is in the native routing instance due to limitations in JunOS. Customer traffic on the MODEL is managed in a virtual routing instance, while customer traffic on the MODEL is in the native routing instance due to architectural differences between the platforms, and the ability to segregate the traffic on the MODEL.

Routing Protocols Definitions

Before discussing the specifics of routing, we must first define some terms:

- Local routes are those configured on the interface on the device and include the logical loopback interface.
- Static routes are configured and are in the form of "get to this prefix via this ip address"
- Learned routes are derived from an active routing protocol.
- Endpoints are the local routes at either end of a physical (logical) interconnect.
- Routing protocols are slippery things, and the same protocol can look internally and externally at the same time.
- Redistribution is when you take routes from someplace, for instance the static table, and add them, or prevent them from being added to another routing protocol.
- Autonomous Systems (ASN) are a collection of routes which may or may not be redistributed.

Even though this is a very small network, we have treated it as a large network from the perspective of routing, so it will be easily extensible and more importantly be maintainable. Best Common Practice (BCP) calls for using three active routing protocols in a network of this type:

- An interior Gateway Protocol (IGP) to carry interface and local routing information. In this case we are use are using OSPF for IPv4 and will use OSPFv3 for IPv6. Although they sound the same, they are not the same protocol.
- An Exterior Gateway Protocol (EGP) to talk to the outside world and link ASN together. The only practical routing protocol to do this is the Border Gateway Protocol (BGP) in its guise, the External Border Gateway Protocol (E-BGP)
- A magic routing protocol to stitch it all together and carry all of the multitude of static routes, external routes, and anything else. In this case we will also use BGP, in its guise of Interior BGP (IBGP).

BGP is designed to efficiently carry large numbers of routes, and to provide fine grained filtering and management which is difficult to do in OSPF.

Routing Protocols Usage

The only routes being carried in OSPF are the local routes, including the loopbacks and interconnect endpoints. This makes for a very small and stable IGP table, since all it's doing is letting BGP know how to get to its next hop. Using OSPF in this way provides us with the ability to do traffic engineering by manipulating routing metrics. In this instance we have two logical connections between Ashburn and DC1, the 10G path provided by the L2 virtual circuit, and a 1G backup path provided by an IPSec tunnel over the commodity internet connections that are used for production and customer traffic. The routing metrics are tailored so all traffic will flow over the 10G link unless it fails, in which case it will transparently move to the backup path. OSPF routes, static routes and local routes are then redistributed into BGP to be carried among all the active nodes in the network. This lets us easily add and subtract prefixes without having to manipulate any routing tables after the initial configuration. The redistribution is also tailored to prevent local routing loops by suppressing the injection of certain routes such as the public IP addresses used to talk to the global internet.

All external communication takes place through the SRX firewalls and is either done over an IPSec tunnels or via a NAT session. In either case, since there is only one external connection at each location a simple BGP default route is injected via EBGP. Using the BGP selection criteria the closest egress point determined by IGP metrics will be used, and if perchance one of the external connections fails, the other connection will be used unless specified by defined routing policy (such as specific NAT definitions and IPSec tunnel endpoints).

Protection features

There are two sets of protections to think about. First, we must be able to prevent traffic from travelling randomly and secondly, we must be as robust as possible.

In order to prevent unwanted traffic in the event of a declared disaster filters will be built to control traffic flow and left in a disabled state where they can be turned on by an operator working from an operations document, which means that a skilled engineer does not need to be contacted in this event.

In the event of a physical failure of equipment which causes an undeclared fail-over (very unlikely given the architecture) when the equipment is replaced and rebuilt from backup, the filter will be activated preventing an unwanted traffic reversion.

Robustness is provided by using segregated redundant paths and high-quality gear.

Backup cluster

[COMPANY] has a backup cluster in DC1 and a standalone backup in DC2. Each site has a licensed capacity of 100 concurrent users and is tied to the Summit backup licensing server.

DC1

- 2 x MODEL Chassis in DC1
- Node0 (-----): id1.servercentral.net
- Node1 (-----): id2.servercentral.net
- Cab: -----

Network Information

- Node0
- External IP: 0.0.0.0
- Internal IP: 0.0.0.0 (VLAN X)
- MGMT IP: 0.0.0.0
- Node1
- External IP: 0.0.0.0
- Internal IP: 0.0.0.0 (VLAN X)
- MGMT IP: 0.0.0.0

Backup Port	Switch Port	Aggregated Interface	VLANs
1	1	-	1, 2
2	2	-	WAN
3	3	-	1, 2
4	4	-	WAB

Cluster Information

The DC1 backup cluster share a floating VIP of 0.0.0.0 so in the event one goes down the other should still be active and accessible via this IP.

Cluster Information

The DC1 backup cluster share a floating VIP of 0.0.0.0 so in the event one goes down the other should still be active and accessible via this IP.

- [COMPANY]_Employees
- [COMPANY] Employees RDP
- Consultants
- RDP Only
- Summit Consultants
- Users - System created Users role

DC2

- 1 x MODEL Chassis in DC2
- Node0 (-----): id1.servercentral.net

Network Information

- Node0
- External IP: 0.0.0.0
- Internal IP: 0.0.0.0 (VLAN X)
- MGMT IP: 0.0.0.0

Backup Port	Switch Port	Aggregated Interface	VLANs
1	1	-	1, 2
2	2	-	WAN

Backups

Due to the limitations of the backup feature on the backups, there is an autoticket in place that drops at the end of every month with instructions on how to take backups. Backups are to be uploaded to location-DC1. From there a cron job will rsync the backups to location-DC1 for redundancy.

Backup instructions are as follows:

- Navigate to Maintenance -> Import/Export -> Import/Export Configuration.
- Under "Export" click the "Save config as..." button to download the System.cfg file. (Even though it says "save as" it will not give you the option to change the file name)
- Switch to the "User accounts" tab and do the same to download the User.cfg file.
- Change the file names of both files to match the following format: user_date_pop.cfg and system_date_pop.cfg - ex. - "user_dc1.cfg".
- Upload the config files to backup@dc1 and their respective directories - /home/----/DC1_backups and /home/----/DC2_backups.
- Leave a note in the ticket that this was completed and close it out.

Firewalls

DC1

- 2 x Firewall Chassis in DC1
- Node0: <location redacted>
- Node1: <location redacted>
- Cab: ----

FW Port Node0	Switch Port	Aggregated Interface	VLANs
1	1	-	1, 2
2	2	-	WAN

FW Port Node1	Switch Port	Aggregated Interface	VLANs
1	1	-	1, 2
2	2	-	WAN

DC2

- 1 x Firewall Chassis in DC1
- Node0: <location redacted>

FW Port Node0	Switch Port	Aggregated Interface	VLANs
1	1	-	1, 2
2	2	-	WAN

Dynamic VPN

At this time the VPN side routes (remote-protected-resources under security / dynamic-vpn) must be manually entered for all internal networks. This means that for every route picked up by OSPF, these must be manually configured under remote-protected-resources.

DC1

Per Summit cluster standard, 1, 2, 3 and 4 are trunked to the top of rack switches (in this case a set of MODEL).

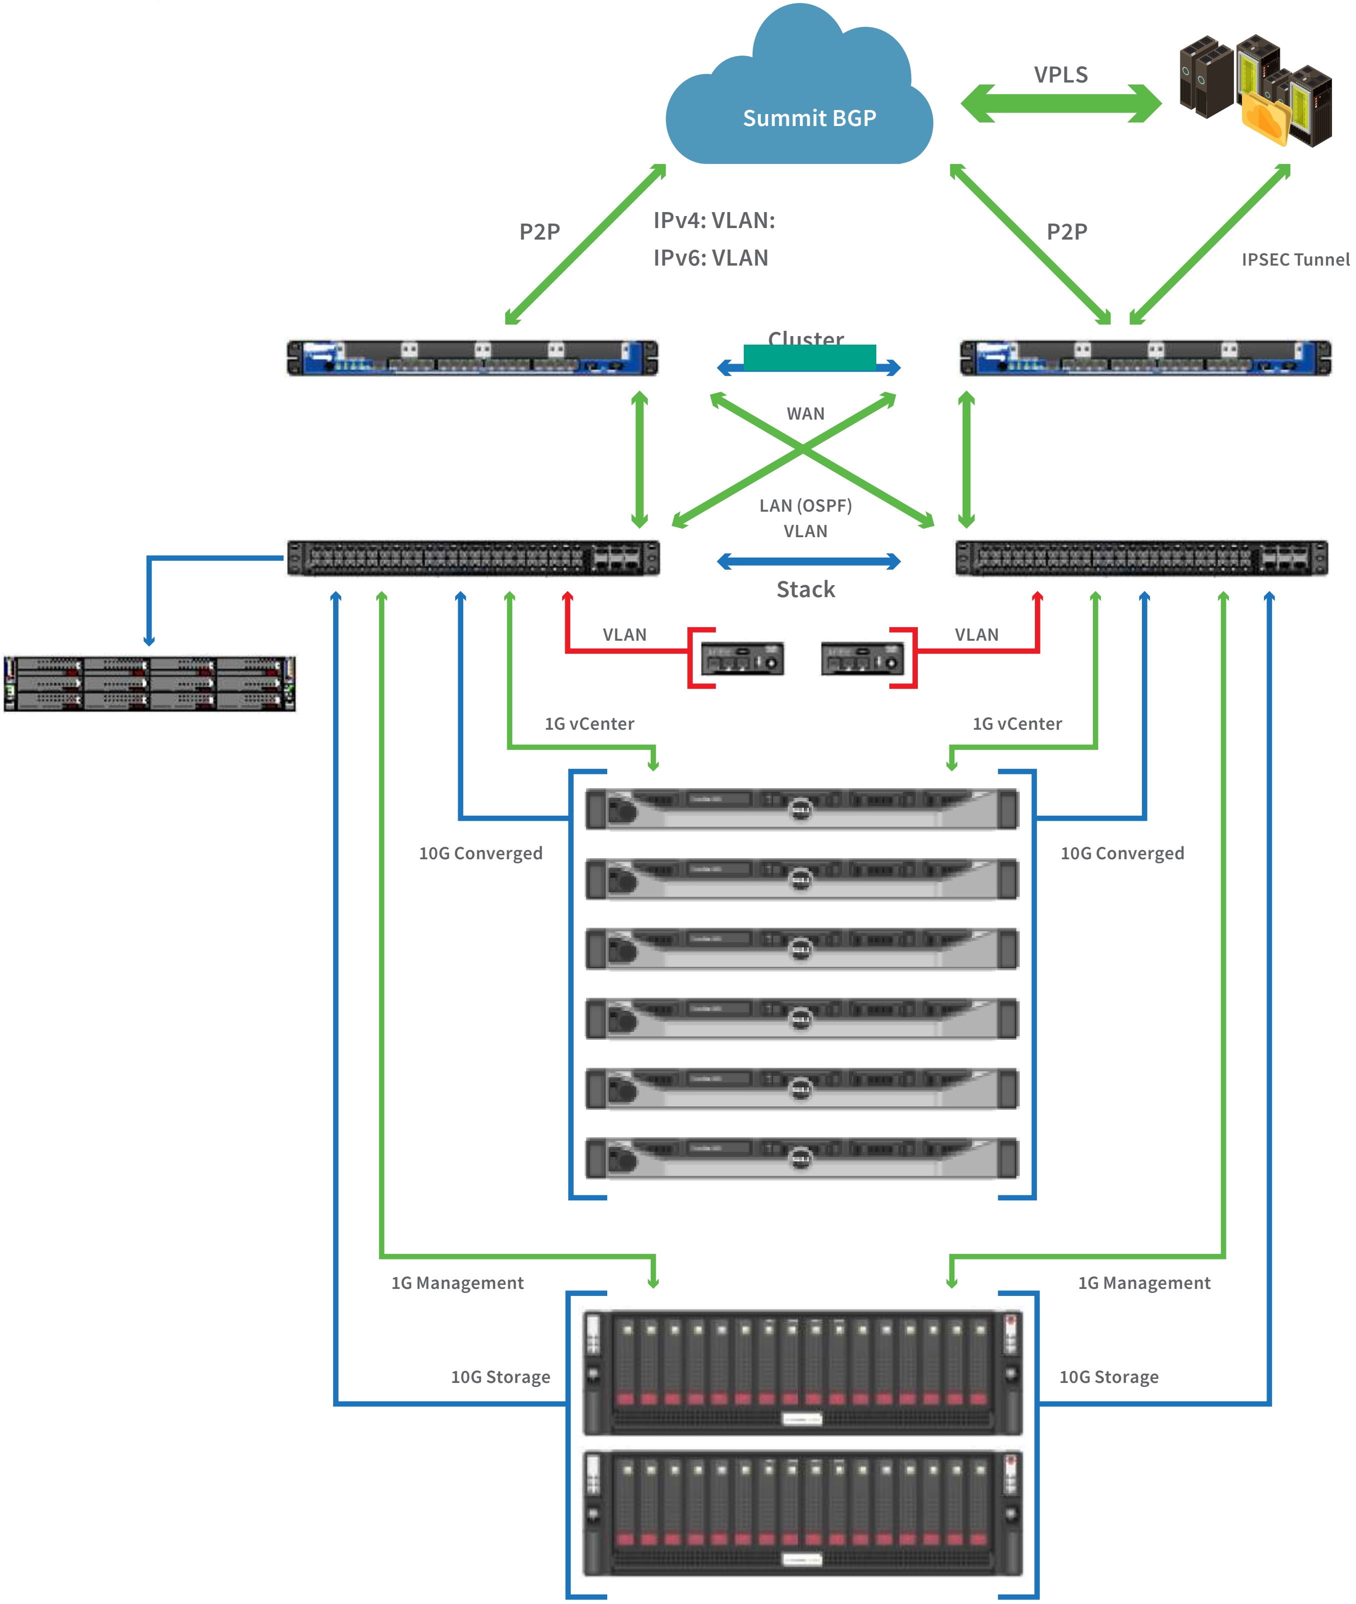

External Side

The [COMPANY] DC1 Firewall cluster is two stacked in LOCATION. VLAN ---- and ---- carry IPv4 point-to-point links to the ARs, and VLAN ---- and ---- carry IPv6 point-to-point links to the ARs as well; these are used for BGP sessions with them. There are policies in place to only import IPv4 and IPv6 default routes from the ARs, and export the current [COMPANY] assigned IP addresses to the ARs.

This /xx sits on a stub VLAN (VLAN ----) in the untrust zone. Dynamic VPN sits on the x.y layer 3 VLAN interface here

Internal Side

The firewall does not sit between internal subnets, and therefore does no internal subnet security policing; it is purely an edge device. Therefore, the firewall has another point-to-point link for obtaining internal routing information and exporting the default route via OSPF. This is 0.0.0.0 on VLAN ----.

OSPF policies are in place to only export a default IPv4 and IPv6 route to the MODEL / [COMPANY] internal network. There is no policy in place for accepting internal network routes, so the SRX imports all of them. Currently the OSPF router-id is 0.0.0.0 hanging off ---- in area 0.0.0.0, but this is causing issues as default outbound IPs will come from this (meaning difficulty with connections to NTP, sysloggers, etc; anything that originates from the firewall). Net Eng is currently mulling a replacement solution; we will most likely just run OSPF directly to an interface rather than the ---- interface.

The [COMPANY] DC1 Firewall cluster is two stacked in LOCATION. VLAN ---- and ---- carry IPv4 point-to-point links to the ARs, and VLAN ---- and ---- carry IPv6 point-to-point links to the ARs as well; these are used for BGP sessions with them. There are policies in place to only import IPv4 and IPv6 default routes from the ARs, and export the current [COMPANY] assigned IP addresses to the ARs.

This /xx sits on a stub VLAN (VLAN ----) in the untrust zone. Dynamic VPN sits on the x.y layer 3 VLAN interface here.

VPN Tunnels

The DC1 firewall cluster has 3 VPN tunnels.

- DC1-DC2 - This is a GRE over IPSEC tunnel that acts as a backup to VPLS between DC1 and DC2
    - Phase1 - Local: ---- Remote: ----
    - Phase2 - Local: ---- Remote: ----
    - GRE - Local: ---- Remote: ----
- This is a policy based IPSEC tunnel
    - Phase1 Local: ---- Remote: ----
    - Phase2 - Local: ---- Remote: ----
- DR - This is a policy-based IPSEC tunnel
    - Phase1 - Local: ---- Remote: ----
    - Phase2 - Local: ---- Remote: ----

DC2

Per Summit cluster standard, 1, 2, 3 and 4 are trunked to the top of rack switches (in this case a set of MODEL).

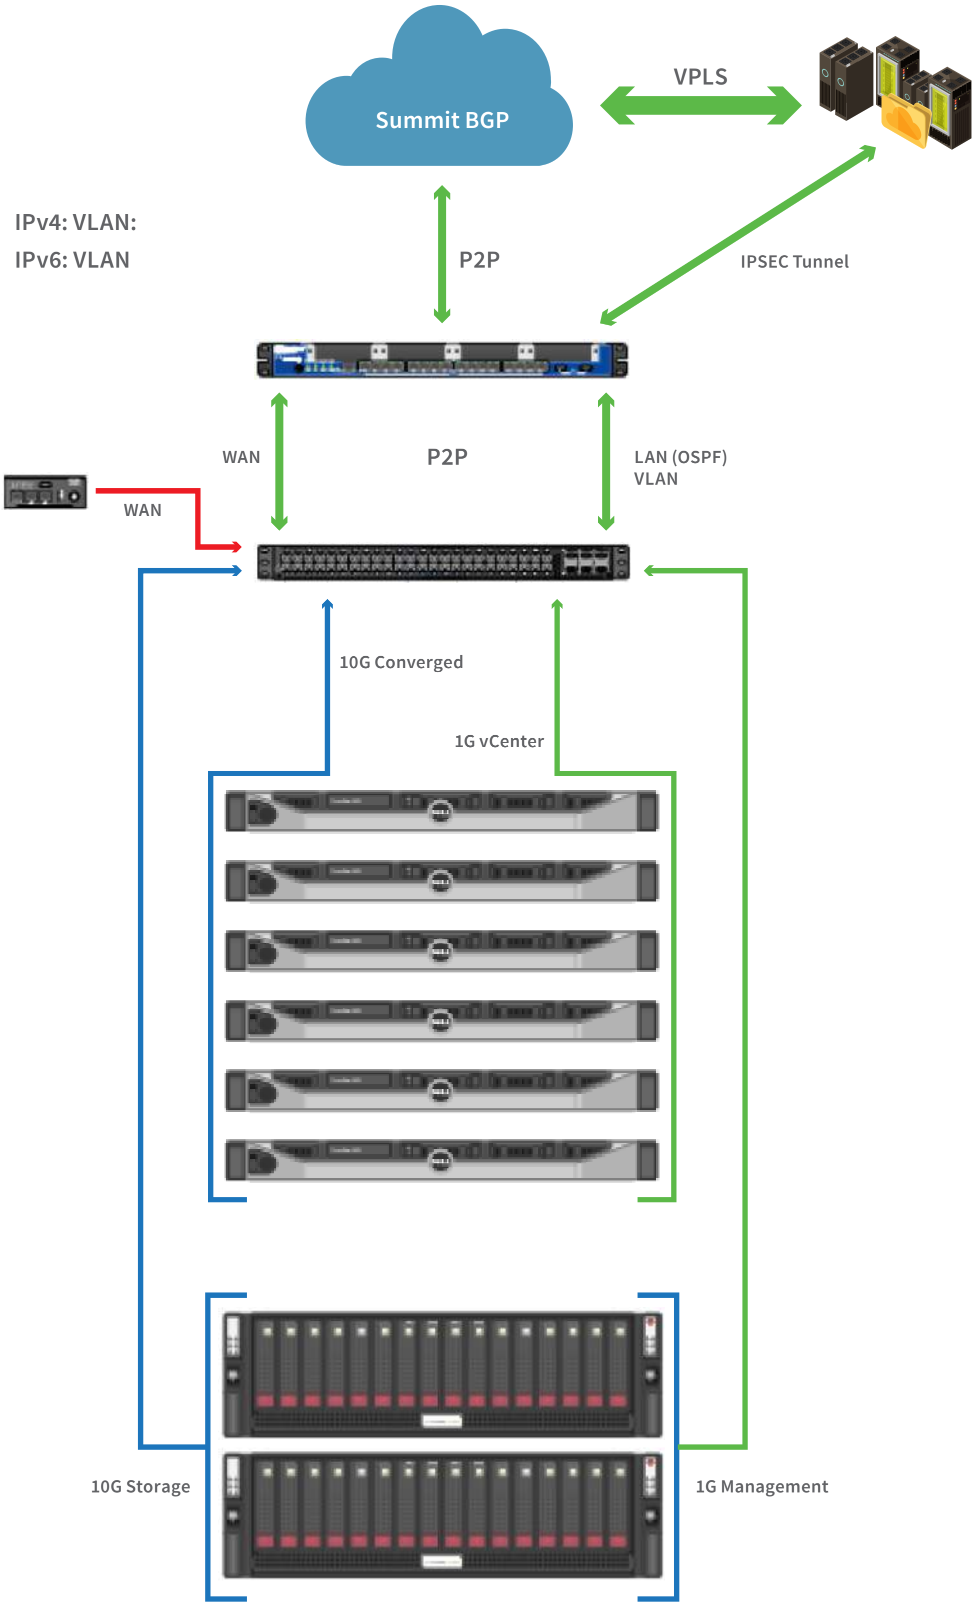

External Side

The [COMPANY] DC2 firewall also carries VLAN ---- and ---- which are unrelated to the DC1 VLANs. These carry BGP sessions to the DC2 AR. VLAN ---- and ---- carry IPv6 point-to-point links to the ARs as well; these are used for BGP sessions with them.

---- is the subnet originating from the firewall. It sits in VLAN ---- .

Note that because of the way routing will work during failover, ---- will be routed back through the QFX, carried over the VPLS circuit back to DC1, and egress there.

Internal Side

The DC2 firewall runs OSPF on ---- @ ----, with router ID ---- on ----.

VPN Tunnels

- DC1-DC2 - This is a GRE over IPSEC tunnel that acts as a backup to VPLS between DC1 and DC2
  - Phase1 - Local: ---- Remote: ----
  - Phase2 - Local: ---- Remote: ----
  - GRE - Local: ---- Remote: ----
- This is a policy based IPSEC tunnel
  - Phase1 Local: ---- Remote: ----
  - Phase2 - Local: ---- Remote: ----
- DR - This is a policy-based IPSEC tunnel
  - Phase1 - Local: ---- Remote: ----
  - Phase2 - Local: ---- Remote: ----

Management IP access – dynamic VPN

In order to access resources within the Managed Infrastructure / Private Cloud enviroment such as vCenter or the SAN admin interface, you will need to connect to the Dynamic VPN using the client in either location first. IP addresses and credentials for DC2 and DC1 are located in:

- Location 1
- Location 2

Please note that you will need to make two connection attempts to the firewall before the connection will succeed (the first connection attempt builds the Phase 1 if its not already established. The second login attempt will complete the connection once the Phase 1 is established). Once connected to either VPN, you can can access both management IP ranges for both PoPs so there's no need to swap VPN connections based on location.

Dyanmic VPN End Points:

- Location 1
- Location 2

Database server cluster

In DC1, [COMPANY] maintains a physical database server cluster. The two machines are 1 and 2. These follow [COMPANY]'s host naming conventions and will not show in vCenter in DC1. These machines have access to communicate with the virtual machines & applications running in the DC1 VMware cluster.

These two nodes are dedicated servers and Summit is responsible for monitoring these servers and maintaining the hardware. [COMPANY] will maintain the operating system & SQL server including any replication.

The DC1 cluster replicates to a virtual machine, DC2, in the DC2 VMWare cluster. There is no dedicated physical server in DC2. It is only the aforementioned virtual machine.

Zerto

- The critical production VMs for [COMPANY] are replicated from DC1 to DC2 using Zerto.
- This is a fully managed service maintained by Summit and related requests are serviced via ticket.
- The source side is DC1 and its replication partner is DC2.
- The passwords for both are kept in the Summit Password Manager.
- Dynamic VPN access into the DC1 and DC2 firewalls is required in order to access this node.

Consult responsibilities

VPG	VM	DR Turn-Up Order
A	vm1	1
B	vm2	2
C	vm3	3

Servers

Server Name	Website	Current IP Internal	DR IP Internal	Current IP External	DR IP External	External Dependency	Turn-Ip Order
A	a.b.c	----	----	----	----	----	9
B	a.b.c	----	----	----	----	----	14
C	a.b.c	----	----	----	----	----	16

Managed Veeam backups

This environment is protected by Managed Veeam Backup. This is a fully managed service maintained by Summit. Backup and restore requests are serviced via ticket.

DC1

The primary backup node is located in DC1. The hostname is ------. The equipment tag is ---- . Credentials are stored in ----. Dynamic VPN access into the DC1 and DC2 firewalls is required in order to access this node. Please refer to DOCUMENT for access details.

DC1-based virtual machines use this target for backups on the following schedule:

Backup Job (Name)	VM	Occurrence	Time	Application Aware	Retention (Days)	DC1 IP	DC2 IP
BU1	A	Daily	----	Y	--	----	----
BU2	B	Daily	----	N	--	----	----
BU2	C	Weekly	----	Y	----	----	----

DC2

The secondary backup node is located in DC2. The hostname is ------. The equipment tag is ----. Credentials are stored in ----. Dynamic VPN access into the DC1 and DC2 firewalls is required in order to access this node. Please refer to DOCUMENT for access details.

Backup Job (Name)	VM	Occurrence	Time	Application Aware	Retention (Days)	DC1 IP	DC2 IP
BU1	A	Daily	----	Y	--	----	----
BU2	B	Daily	----	N	--	----	----
BU2	C	Weekly	----	Y	----	----	----

Summit virtual machines

Summit operates a number of "utility" virtual machines that are necessary for our management of the infrastructure. Dynamic VPN access into the DC1 or DC2 firewalls is required in order to access this node. Please refer to DOCUMENT for access details.

DC1

- vCenter Server Appliance
- VMware Update Manager
- Zerto Virtual Manager
- vCloud Usage Meter
- LogicMonitor Collector

DC2

- vCenter Server Appliance
- VMware Update Manager
- Zerto Virtual Manager
- vCloud Usage Meter
- LogicMonitor Collector
- Veeam Backup Proxy

Data center operations

LogicMonitor Notifications

Send all [COMPANY] alerts to servermonitoralerts@company.com as outlined in DOCUMENT.

All other Summit infrastructure or Managed Services alerts should be escalated per alert text or Summit escalation procedures.

- End of Managed Infrastructure / Private Cloud Section -

Disaster recovery as a service (DRaaS)

This section shall establish the Disaster Recovery runbook for the [COMPANY] managed infrastructure / private cloud environment hosted with Summit. This is a step by step, actionable plan with clear indications for what is required from [COMPANY] and Summit during each portion of the fail-over and fail-back procedures.

Disaster definition

[COMPANY] has defined a "disaster" into two terms:

Provider Disaster

- Defined as an outage of the infrastructure provided by Summit.
- During a Provider Disaster, Summit will notify [COMPANY] of the outage. However, [COMPANY] must still authorize fail-over procedures to be executed and to officially declare a disaster. Summit cannot declare the disaster themselves nor execute the fail-over procedure without explicit authorization from [COMPANY].
- [COMPANY] will allow Summit up to two (2) hours to provide an estimate for outage resolution.
- After three (3) hours elapsed, [COMPANY] will notify Summit to prepare for a disaster declaration and for Summit to gather the technical resources necessary to execute the fail-over procedures.
- If the time to resolve the outage is >4 hours, disaster will then be declared by [COMPANY].
- Disaster may be declared at any time prior to the 4-hour elapsed mark.

Application / Software Disaster

- Defined as an outage related to Application or Software failures within the COMPANY] hosted application.
- During an Application/Software Disaster, [COMPANY] will troubleshoot the application issues for up to 4 hours before declaring a disaster.
- After three (3) hours elapsed, [COMPANY] will notify Summit to prepare for a disaster declaration and for Summit to gather the technical resources necessary to execute the fail-over procedures.
- If the time to resolve the outage is >4 hours, disaster will then be declared by [COMPANY].
- Disaster may be declared at any time prior to the 4-hour elapsed mark.

Disaster recovery administrators

Only [COMPANY] personnel with the DR Admin flag (denoted as a D within Summit's inventory & portal system) can:

- Declare a disaster to initiate the fail-over procedure.
- Declare a disaster over and initiate the fail-back procedure.
- Request DR/sand box testing.

Before executing any procedures, Summit personnel will verify identity with the contact information on file. The most current list of DR administrators will appear within the Summit Customer Portal/ACL system.

[COMPANY] disaster recovery team

[COMPANY] has identified the following internal teams as responsible for being present during a DR event. This list is maintained by [COMPANY] and [COMPANY] must notify Summit of any changes to this list. These individuals will be joining the DR Conference Bridge during an event.

Incident / Infrastructure

- First Last / First Last

Executive

- First Last / First Last

AppDev

- First Last / First Last

Data Team

- First Last / First Last / First Last

Disaster recovery conference bridge

[Company] has supplied a conference bridge for all parties to join during a declared disaster event. After verifying the identity of the DR admin declaring a disaster but prior to executing any procedures, Summit Managed Services will join the DR conference bridge to coordinate recovery efforts with [COMPANY]. This bridge has been confirmed to be powered by an independent phone system that is separate from the main [COMPANY] phone system and will be available during a [COMPANY] phone system outage.

Conference Bridge Information

- Toll Free Dial-In Number (US & Canada): ----------
- US local toll Dial-In Number: ----------
- Conference code: ----------
- Leader PIN: ----------

Disaster recovery prerequisites

[COMPANY] is responsible for confirming the following:

- Confirmed SSL VPN functionality in DC2.
- Confirmed working directory services in DC2.
- Confirmed working database replication from DC1 to DC2.
- Confirmed Zerto VPG turn-up priority.
- Provide backup connectivity from LOCATION to DC1.

Summit is responsible for confirming the following:

- Confirmed site to site tunnels are online and passing traffic.

Fail-over procedure

Fail-Over – Identity Verification

An authorized [COMPANY] DR Admin shall submit a ticket request to Summit declaring disaster and to execute the fail-over procedure.

- Summit will verify the identity of the ticket submitter by verifying the email address of the ticket submitter matches what is on file in the Summit ACL.
- Summit will contact a second DR admin contact to confirm the disaster.
- Summit will then inform the NOC of the active disaster.

Fail-Over – Join DR Conference Bridge

- Summit Managed Services will join the [COMPANY] DR Conference Bridge. See DOCUMENT for details.
- [COMPANY] will provide executive confirmation of a disaster being declared.
- [COMPANY] DR team and other needed staff will join the DR Conference bridge.
- Once all necessary personnel are gathered, [COMPANY] & Summit will then begin executing the appropriate procedures.

Fail-Over – Phone System

[COMPANY] Responsibility:

- Execute phone system fail-over procedure from DC1 to DC2.
- Contact PROVIDER customer service at ----------.
- Select option X, option X, option X, option X
- Provide the [COMPANY] account number: ---- ---- ----
- Inform agent to disable trunk group ---- (DC1).
  - NOTE: Disabling trunk group ---- will force calls to go to trunk group ---- (DC2).
- Confirm phone system operation.
  - Make a call to (571) 206-3452
  - Verify the call goes through to the local time server for the US Naval Master Clock.

Fail-Over - Power Down Dev/QA VMs – DC2

Summit Responsibility:

- Power off Dev/QA virtual machines for all machines marked as Dev/QA = Yes. See DOCUMENT for a list of VMs to power down.

Fail-Over – Database Cluster

[COMPANY] responsibility:

- In the event that the Production Database Server Availability Group (------) needs to become Primary in DC2 and there is communication between the two datacenters:
- Sign onto DC2-----.
- Open the Failover Cluster Manager and make DC2----- the Primary in the Cluster.
- On personal computer, open Database Server Management Studio, connect to DC2----- and run:
  - Step 1
  - Step 2
- Verify database functionality and operation.

Fail-Over – Zerto – Critical VMs

Summit Responsibility

- Verify a split-brain situation is not present between DC1 and DC2 VMs.
    - If a split-brain situation exists, power off all DC1 VMs.
- Execute Zerto VPG fail-over based on turn-up priority listed in DOCUMENT
- Verify Guest OS booted to login prompt via VMware Console.
- Confirm DR IP addresses are reported in vCenter on a per VM basis (VMWare tools). See DOCUMENT.

[COMPANY] Responsibility

- Execute the internal DNS update script.
- The scripts can be located here: FILE LOCATION
- NOTE: The script must be run from a workstation with remote server admin tools installed. A domain controller can be used if a workstation is not available. The script requires the DNS powershell cmdlets.
    - Start a Windows Powershell console as an Administrator by right clicking the Powershell icon and selecting "Run as administrator". Alternatively, you can also click Start > Run > then typing “powershell” into the dialog box then right clicking on "powershell.exe" and clicking "Run as administrator".
    - Once the PowerShell console opens, type "SCRIPTNAME" and press Enter.
    - The script will execute. While the script is running, the statements in yellow text will confirm the DNS entry was updated successfully:

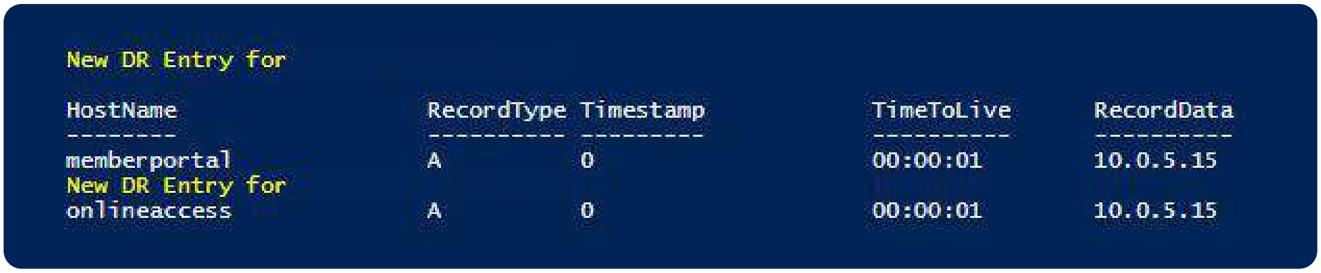

```
New DR Entry for

HostName                    RecordType Timestamp             TimeToLive       RecordData
--------                    ---------- ---------             ----------       ----------
memberportal                A          0                     00:00:01         10.0.5.15
New DR Entry for
onlineaccess                A          0                     00:00:01         10.0.5.15
```

- Verify ability to login to Guest OS.
- Confirm and fix proper IP address assignments to the guest OS.
- Execute IIS re-binding scripts on affected VMs.
- Confirm site to site tunnels operational from [COMPANY] HQ and LOCATION.
- Perform remaining fail-over steps per the following table of operations and responsibilities:

Duration (Minutes)	Task	DR Team	Notes
10	Verify server replication was successful	First Last First Last	---
15	Verify directory services and DNS are running and functional	First Last First Last	---
5	Recovery servers volume by volume	First Last First Last	---

Fail-Over - Apply Switch Filters - DC1

Summit Responsibility:

- After Zerto - Critical VM Fail-Over is completed, if DC1 is still accessible, apply switch filters in DC1 to prevent site to site split brain communications between the now production VMs in DC2. LOCATION will still be able to access DC1 for troubleshooting purposes once DC1 connectivity is restored.
    - On switch, activate filter term ---- using normal Juniper commands

Fail-Over - Dynamic DNS Updates

During a DR event, external DNS will need to be updated. [COMPANY] uses PROVIDER for their external dynamic DNS provider and has created a script to push the necessary external DNS changes via the API.

- The scripts can be downloaded here: FILE LOCATION

Customer Responsibility:

- Execute DNS PowerShell update script to re-point customer websites to DC2 public IP addresses.
- Step-by-Step instructions:
    - Start a Windows Powershell console by clicking the Powershell icon or by clicking Start > Run > then typing “powershell” into the dialog box and pressing Enter.
    - Once the PowerShell console opens, type "SCRIPT NAME" and press Enter.
    - The script will then execute and prompt for credentials for DNS authentication. The credentials can be located in LOCATION
    - After the credentials are supplied, the script will execute. A "good" status will be displayed along with the IP address of the DNS entries that were successfully modified:

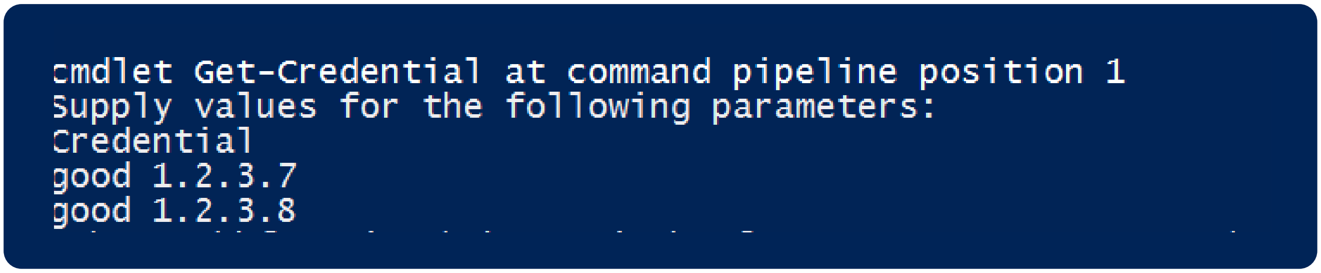

```
cmdlet Get-Credential at command pipeline position 1
Supply values for the following parameters:
Credential
good 1.2.3.7
good 1.2.3.8
```

Fail-Back - Clean-Up Activities

- Re-establish original Zerto virtual protection groups from DC1 to DC2.
  - NOTE: This must be re-established using the original IP addresses.
- Re-establish original Veeam backup schedule in DC1 and DC2.
- Re-establish original monitoring protocols.
  - Inform NOC to resume normal monitoring procedures.

Disaster recovery sandbox environment

Purpose

We have created a completely segregated network in order to test failover of production VM's and to test functionality between various VM's and applications. Once placed into the proper sandbox networks, they can only communicate with other sandbox networks and nothing else.

Requirements

Due to the nature of network segregation, the VMs need access to certain VMs for AD services and SQL data. These VMs are: 1 / 2 / 3

[COMPANY] accesses the sandbox environment via these VMs: 1 / 2 / 3 / 4 / 5

The sandbox networks: 1 / 2 / 3 / 4 / 5

Clone local resources

Due to the size of the VMs listed in these requirements, it is quicker to deploy from a SAN clone than it is to make a VMware clone.

- Log into the DC2 vCenter via the web client.
- Verify via the VM settings that all the required VMs are on the same datastore
- Use the Nimble Storage Plugin to clone the datastore
- Name the new datastore and create a new snapshot
- Once the clone process is complete. Browse the newly added datastore for the required VM's and register them.
- Make sure to append -sandox to the VM name and to place it in the Failover-Testing Folder.
- Select the main the cluster.
- Edit the VM settings to change the NIC to the proper Sandbox network before powering on. The VM may have an alert about a duplicate MAC address but this is expected.

Once the VMs have booted up, you can hand off verification to customer before progressing into production VM failover.

Zerto Test Failover Procedure

Typically, we have been failing over VM's in groups of 5. Customer will tell you when to progress to next group.

- Log into the DC1 ZVM and select the VPG tab.
- Select the proper VPG's based on the proper order listed on this page.
- Click on the large Failover button at the bottom (default to Test).
- Click Next. Click Next. Click Start Failover Test
- Zerto will then create and start the selected VM's in the Sandbox. Once the VMs have booted and VMtools are up and running. It will execute scripts to set the IP's and change the network settings.
- Once the VM's have rebooted, verify they are displaying the correct IP in vCenter and notify customer they are ready for verification.
- Repeat as necessary.

Sandbox Testing Cleanup

Once all testing is complete, we will get the all clear to cleanup the environment.

- From the DC1 ZVM, select the VPG tab. All the failed over VPG's will have an operation value of "Testing Failover". For each of those, click on the red square to Stop Failover Test.
- From the DC2 vCenter, power off the local required VM's in the sandbox.
- Remove the same from Inventory.
- From the datastore view, verify there are no VM's listed on the sandbox-test datastore
- Use the Storage Plugin to delete the datastore.
- Once removed, you can safely remove the created snapshot from the Primary4 datastore.

- End of Document -