ZYNSTRA TECHNICAL BRIEFING NOTE

Network Integration Considerations

Network Connections

Each Cloud Managed Server has at least five built-in Ethernet ports which are used as follows:

- A **Wide Area Network** (WAN) port which connects the server to the Internet.
- A **Local Area Network** (LAN) port which connects the server to the Customer's LAN and which is only ever permanently disconnected if the Customer's LAN is not in the same building as the server itself but is reached over the WAN connection.
- A **Remote Management** (iLO) port based on the HP iLO (integrated Lights Out) capability which enables access to the server after a serious failure, even if the WAN connection is unusable.
- Two **Cluster Interconnect** ports (CL1, CL2) which are only used in the Customer has purchased a High Availability (HA) cluster and which directly connect the two Cloud Managed Servers in the cluster to each other.
- The example below shows the various ports and their usage on an appliance that has been sold with the HA option (and hence comprises two servers in a clustered configuration).

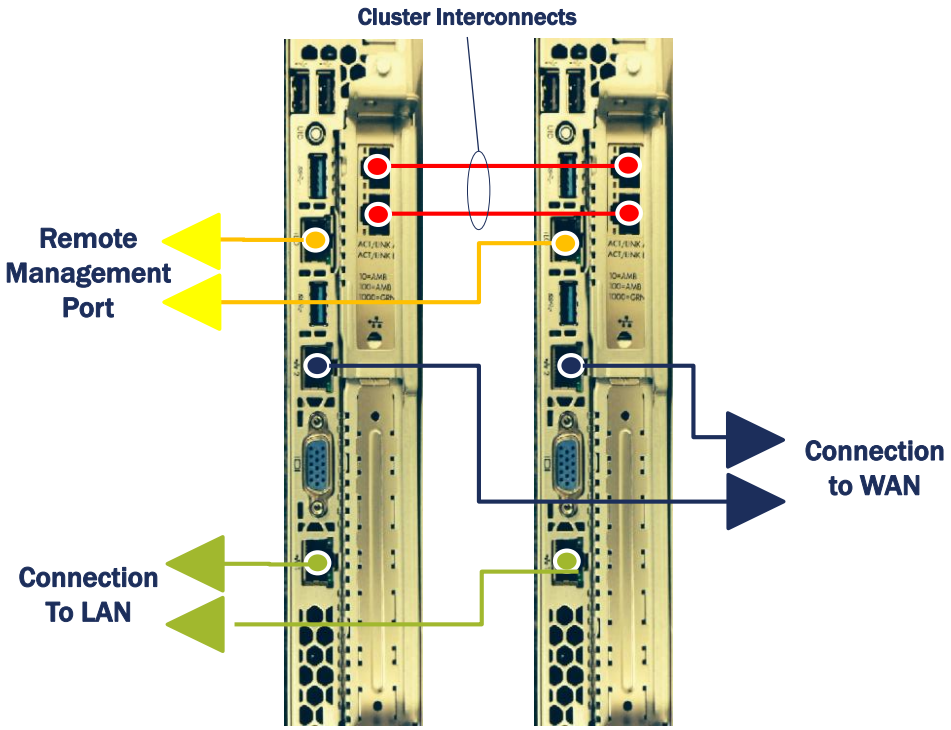

IP Address Requirements

For a standard Cloud Managed Server comprising a single node, the following IP addresses are required:

- **WAN IP address**: **either** a public IP address to directly connect the server to the Internet through the Customer's router **or** a private IP address to connect the server to the Internet through the Customer's firewall
- **LAN IP address**: a private IP address to identify the server to the devices on the Customer's LAN
- **iLO IP address**: **either** a private IP address for the remote management connection to enable the iLO to be permanently connected to the Internet through a firewall or a public IP address to enable the iLO to be connected to the Internet on-demand and only when remote diagnosis is required.

For servers to be deployed as a High Availability solution comprising two servers acting as a cluster, the following IP addresses are required:

- **One WAN IP address**: **either** a single public IP address to directly connect the active node in the cluster to the Internet through the Customer's router **or** a single private IP address to connect the active node in the cluster to the Internet through the Customer's firewall
- **One LAN IP address**: a private IP address to identify the active node in the cluster to the devices on the Customer's LAN
- **Two iLO IP addresses**: **either** two private IP addresses to enable the iLO to be permanently connected to the Internet through a Customer firewall or two public IP addresses to enable the iLO to be connected to the Internet on-demand and only when remote diagnosis is required.

For the High Availability solution, it is important to note that only one WAN IP address and one LAN IP address are needed even though there are two servers as the WAN and LAN IP addresses are only active on one of the two servers at any one time and, in the event of a failure of the server on which they are active, they will move automatically to the other server.

All public IP addresses used by the server must be routing IPv4 addresses with a router default gateway in the same public subnet. They should not be bridged, dynamic or PPPoE/PPPoA and should have no NAT, QoS, inbound or outbound firewall rules or other traffic management.

All the IP addresses above must be static IPv4 addresses.

Some applications and services running within an appliance may themselves require static public IP addresses in order to operate correctly.

Most Common Network Topology Before Introducing the Cloud Managed Server

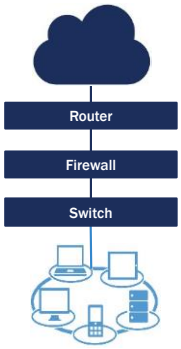

The most common arrangement of a Customer's LAN and its connection to the Internet prior to the introduction of the Cloud Managed Server is as shown to the left. The Customer has a router, often the device supplied by their Internet connection provider. They have one or more static, public IPs available to them from the provider of their Internet connection, at least one being used for this WAN connection from the router. They have a firewall and sometimes the firewall and the router are combined into a single device. The firewall connects to the devices on the LAN through a LAN switch.

'Replace Existing Firewall' Scenario

The most common deployment scenario is one in which the Cloud Managed Server replaces the customer's existing firewall and therefore connects on the WAN side into the Customer's existing router (which must present an Ethernet interface) and on the LAN side into a LAN switch. Two static, **public** IPv4 addresses are required for this scenario (WAN IP, iLO IP) or three if it is to be a High Availability solution (WAN IP, iLO IP server 1, iLO IP server 2) so a key pre-requisite is that the customer has or can acquire these additional static, public IP addresses. This scenario cannot be used if the Customer's firewall and router are combined in a single device and the customer wishes to retain this device unless the combined device can be reconfigured by the customer to act purely as a router or the combined firewall and router can be replaced with a pure router. One private LAN IP address is required which can be the one used already for the LAN side of the firewall.

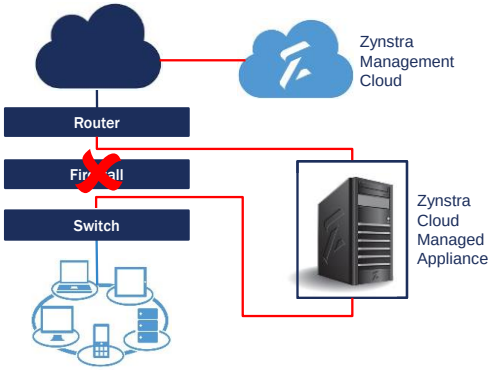

'Behind Existing Firewall' Scenario

If the customer's router and firewall are combined in one device which they wish to retain, or if the customer can only source one public IP address, or if there are other reasons why the customer wishes to place the appliance behind their existing firewall, this scenario can be used. Two static, private IP addresses are required (or three for the HA option) and a double-NAT is required. One private LAN IP address is required which can generally be the one used for the LAN side of the existing firewall.

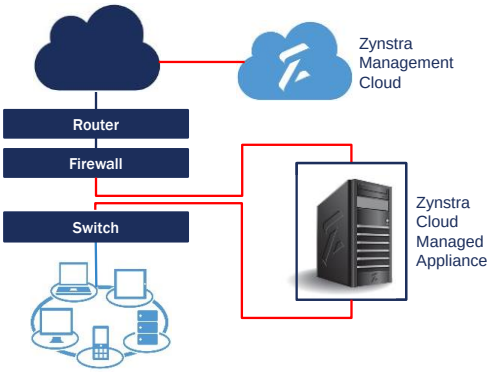

‘Parallel To Existing Firewall’ Scenario

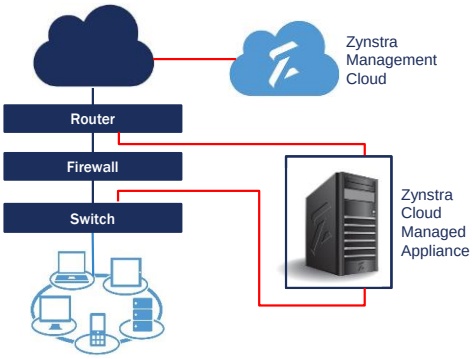

If the customer wishes to retain its existing firewall because it is serving other purposes beyond the connection to the LAN switch, and the customer can supply the additional public IP addresses required, the appliance can be deployed parallel to the existing firewall and connect directly into the router (provided that it presents an Ethernet interface and has sufficient spare ports). Two static, **public** IP addresses are required for this (or three if the HA option has been purchased) so a key pre-requisite is that the customer has or can acquire these additional static, public IP addresses. One private LAN IP address is required which can be the one used already for the LAN side of the firewall.

‘Third Party Data Centre’ Scenario

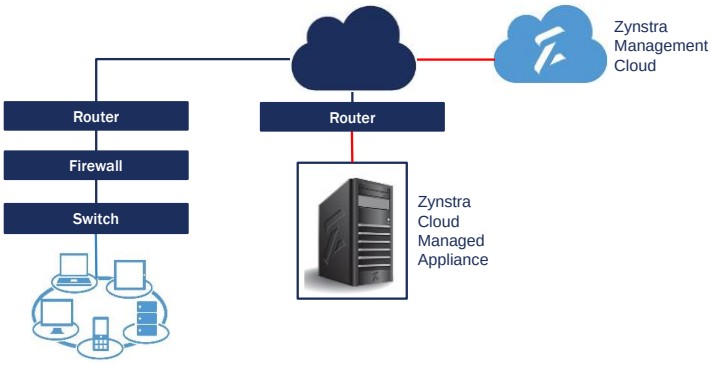

If the customer wishes to locate the appliance in a third party data centre, this can be achieved by placing it behind a router provided that the router presents an Ethernet interface and that the required public IP addresses can be supplied.

Two static, **public** IP addresses are required (or three with the HA option). No LAN IP is required for the appliance as users will access the capabilities of the appliance over the Internet through a VPN connection from their personal devices.

Other Network Integration Scenarios

Other scenarios (as well as permutations of the above scenarios) are possible when integrating a Zynstra appliance into a customer’s network. It is recommended that any customer requirements which do not fit exactly with the above scenarios are discussed directly with Zynstra at the earliest possible opportunity to verify their feasibility, including customer networks which currently comprise use VLANs to create multiple layer 3 networks on a single physical infrastructure.

If a customer has multiple sites, and if the customer’s existing firewall is terminating one or more site-to-site IPSec VPNs or any client-to-site VPNs, further discussion is required with Zynstra to assess the most appropriate approach to integration unless the appliance is going to replace the existing gateways on all the customer’s sites at the same time.