Best Practices

Best Practices in Application Configuration

We would like to offer some best practices in setting up redundant, available, and efficient applications. This list will grow, but for now, we offer the core best practice: basic redundancy of important instances. **Nothing can prevent random equipment failure, the occasional malfunction of an ISP, or even the rare catastrophic event.** For this reason, our SLA is based on the kind of best practices redundant setup that makes sense in any datacenter.

Basic Redundancy: At Least Two Copies of an Instance

Because of random hardware failure, it always makes sense to run two instance copies on two different servers. But global redundancy is an even better way to prevent downtime; then if an entire datacenter is affected, your applications stay live. AppNexus currently offers three datacenters: one in the New York region, one in Los Angeles and one in Amsterdam. For every key instance you run, we strongly advise setting up a load-balanced copy in the LAX1 datacenter, one in the NYM1 datacenter, and one in the AMS1 data center if your business include Europe.

Note that if your applications do not require the resources of an entire server, two or more separate applications could be made fully redundant using a total of two servers.

Steps to create two redundant instances:

1. Make a copy of your instance. See Bundle an Appnexus Instance for details.
2. If you currently have a VLAN in only one datacenter, please create a ticket at https://portal.appnexus.com/ or contact us at support@appnexus.com to be assigned a VLAN in the second datacenter.
3. Use `rsync` to copy your instance to the second datacenter. Launch copies of your instances in each datacenter. See Start an Instance from a Custom Image for details.
4. Contact AppNexus Support to set up global load balancing between the two instances. See Managing Global Server Load Balancing for details.
5. It is also a good idea to save a copy of any key instances on Network Attached Storage as an added backup.

Further Reading

Load Balancing Overview
Configuring Local Load Balancing
Managing Global Server Load Balancing

Using Jump Boxes to Limit Access to your VLAN

For security, we recommend locking down your entire VLAN and only opening really necessary ports and source and destination IP addresses. To access your VLAN, you would set up two jump instances on two different host servers and use these boxes for SSH connection to all your other instances.

Further Reading
Other Security Recommendations

Using the Load Balancer

It makes sense to run all externally facing services--~~even single applications~~--through the local traffic manager (LTM) for several reasons:

1. Load-balancing pools make it easy to migrate applications instantly. With a single node, the DNS Time to Live (TTL) will have to expire before your users see the backend change you have made. When user traffic is passed through an LTM, the LTM will redirect instantly, as soon as a change is made.
2. It is easy to add servers and capacity in an instant. Your pool is already configured, and adding a node takes a single command. (`manage-lb-pool add-node`)
3. The LTMs have built-in protection against Distributed Denial of Service (DDoS) attacks. This adds a layer of security.
4. Direct connections to backend servers are prevented, for security reasons.

Monitoring Your Instances

Another best practice is monitoring your instances for information on system load and connectivity issues. We recommend the Ganglia tool for this. For instructions, please see Monitoring Instances Using Ganglia. (For information on the monitoring that AppNexus does on core infrastructure, see here.)

As always, please create a ticket at https://portal.appnexus.com/ or contact us at support@appnexus.com if you have any questions or concerns.