Infoblox Grid™ Technology

Enabling Next-Generation Reliability for Core Network Services

Executive Summary

Infoblox delivers network core network services—including DNS, DHCP, and IP address management (DDI)—to on-premises, cloud and hybrid networks. The Infoblox Grid™ is Infoblox's unique and patented high-availability (HA) technology for ensuring the reliability of these services. It provides core network service resiliency, failover, recovery, and seamless maintenance for an Infoblox deployment inside a single building, across a networked campus, or between remote locations.

An Infoblox Grid is created by linking physical and virtual DDI appliances together using a secure communication platform. The solution is not a separate management and reporting application that overlays individual appliances. Rather, appliances in an Infoblox Grid are linked to support a sophisticated distributed database embedded within each appliance. This transforms a set of appliances into a unified and resilient system that removes single points of failure and other operational risks inherent in legacy DNS, DHCP, and IP address management infrastructure.

The Infoblox Grid addresses the basic problems that remain unresolved when independent servers or appliances are deployed within a distributed enterprise. For example, in such implementations, each server or appliance must be individually deployed, managed, and upgraded. Because each server or appliance acts on its own, it cannot ensure the availability, accuracy and timeliness of network service data. The end result is that these individual appliances cannot act as an integrated system, cannot offer high availability, are not robust in the face of network outages and increase the burden on IT staff.

The Infoblox Grid address all of these issues, and more.

Core Network Services—Key to All IP Network Devices and Applications

Core network services deliver and manage critical data that maintain the relationships between users, their IP-connected devices, user privileges, and network resources. This critical data includes:

- IP addresses
- Domain Name System (DNS) naming services
- Domain Name System Security Extensions (DNSSEC)
- IP addressing services via Dynamic Host Configuration Protocol (DHCP)
- Network visibility and control via IP address management (IPAM)

Device proliferation is driving an exponential increase in demand for core network services. Along with the explosive growth of mobile devices, the rise of the Internet of Things (IoT) has resulted in a multitude of "smart devices" clamoring for their own place at the table for network services. The vast array of smart devices includes everything from ATMs to cameras and wearables to connected manufacturing equipment, connected vehicles, and smart appliances. Each device requires its own IP address. Other services that draw on core network services including email, web services, voice over IP (VoIP), Microsoft Active Directory, and wireless networking. All of these concurrent demands, from both devices and applications, dramatically increase the scale and scope of delivering network services. Ensuring the availability and simplifying management of these services is a top IT priority.

Infoblox Grids Raise Reliability to a New Level in Distributed Environments

The Infoblox Grid enables an organization to distribute and consolidate critical information and services in real time with assured data integrity, including:

- Protocols (DNS, DHCP, TFTP, NTP, etc.)
- Data (IP addresses, MAC addresses, user credentials, transaction logs, time, etc.)
- Files (firmware images, configuration files, policies, etc.)

The need for this type of infrastructure is especially acute in light of trends in IP networking, as discussed below.

The Need to Distribute as well as Consolidate

A number of key IT imperatives are driving a need to distribute core network services, while others are driving the need to consolidate management and control for these services through a single pane of glass.

Needs that compel an enterprise to utilize distributed core network services include:

- Application performance: Server and propagation delays between clients and DNS servers can be the limiting factor in application performance, especially as browser-based and cloud-based applications increasingly draw static and dynamic content from different servers that require multiple DNS requests per page. Without survivable, high-performance local DNS services, branch applications may exhibit poor performance or cease to function.
- Disaster recovery: In the event a key server site is lost or experiences connectivity issues, core network services may need to be sourced from geographically dispersed locations.
- Local survivability for VoIP: Without highly available DHCP and TFTP services to deliver IP addresses and firmware images to IP-based phones, remote users may be unable to make or receive phone calls.

At the same time, the following forces are driving a need for consolidating network services data and management:

- IP address management (IPAM): as noted above, the rapid growth in the number of IP-based devices, including phones, tablets, RFID systems, and others—in addition to PCs and servers—is driving a need for IPAM solutions that enable real-time allocation and tracking of IP addresses assigned to the devices.
- Regulatory compliance: Sarbanes-Oxley and other regulations demand centralized management and reporting of administrative changes and network activity across the enterprise.
- Resource constraints: The limited availability of skilled IT personnel and demands for increased operational efficiency require reduced administrative overhead for deploying and managing core network services. This drives the need to delegate administrative responsibility for low-level repetitive tasks while enabling unified, system-wide visibility and control for senior IT staff.

The Infoblox Grid is unique in its ability to resolve these conflicting requirements, providing highly available and secure local service delivery with the benefits of unified management and control.

Infoblox Grid Technology Overview

Infoblox Grid technology enables distributed Infoblox appliances to function as a unified, centrally managed system—instead of independent devices. This capability provides a real-time distribution, synchronization, and management framework. The Infoblox Grid is implemented using appliances licensed with the NS1-Grid package, which enables the functions of the Grid module in the Infoblox Network Identity Operating System (NIOS™) software. The Grid module leverages and enhances underlying subsystems in the Infoblox NIOS software included with every Infoblox appliance.

The NIOS software includes service modules such as DS, DNSSEC, DHCP, TFTP, and NTP services that are implemented using industry-standard protocol engines and a database engine. These engines provide:

- Zero-administration: The database is built-in and requires no user installation or maintenance for database replication and distribution.
- Persistent transactional subsystem: Ensures no data loss throughout a single or distributed system, even in the event of a failure.
- Semantic constraints: Provides data validation and consistency checks.

It is important to note that the database is able to provide these services to all protocols supported in the NIOS software. This makes it possible, for example, to transform a protocol module such as TFTP, which has no inherent concept of distributed operation across multiple systems, into a Grid-enabled protocol that provides central management and ensures file consistency across sites.

Floating Master Architecture

In an Infoblox Grid, at least one of appliance is designated as the Grid Master and is responsible for coordinating and synchronizing data and configurations across the other appliances, which are designated as Grid Members. Grid Members serve local DNS, DNSSEC, DHCP, TFTP, and NTP data (via proxy). The Grid Master may also serve data, but it also has several special roles:

1. Provides the seat of administration for the Grid: The Infoblox Grid Manager application communicates with the Grid Master, which in turn provides configuration data to each of the member appliances in the Grid. Figure 1, featured below, illustrates data visibility among Grid members.

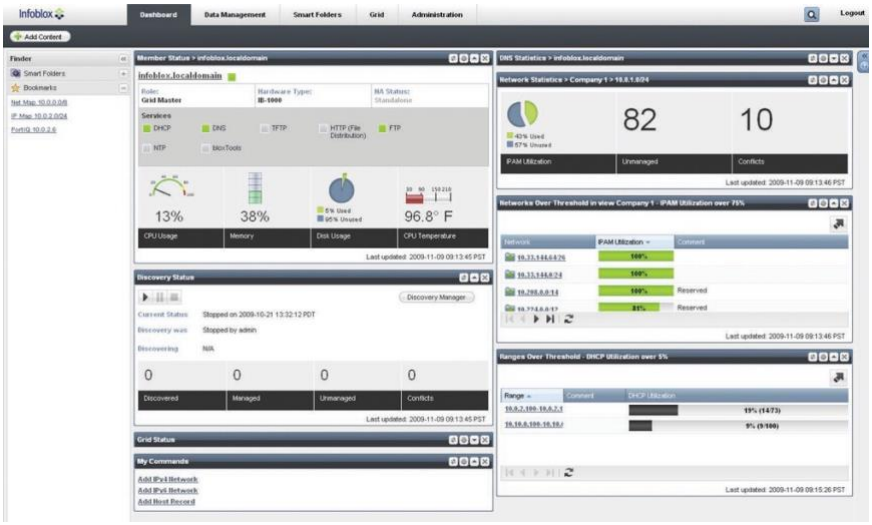

Figure 1: The Infoblox Grid Manager provides a unified view of an Infoblox Grid

2. Manages software updates: New software revisions for appliances in the Grid are uploaded to the Master. The Master is responsible for managing the software distribution and upgrade process on remote members. Files for delivery via TFTP are also distributed to member appliances via the Master.
3. Synchronizes Grid data: Changes that occur on member devices are transmitted in real time from each member to the Grid Master. (A change to all DNS zone data due to dynamic DNS updates is one example of the type of change that would be transmitted in real time.)

   Once this data has been received by the Grid Master, it sends these updates to Grid member devices. The Grid Master knows which member devices serve the updated data and transmits data only to those devices affected by the update.

   Data is partitioned so that member appliances contain only data that they serve. The replication mechanism is intelligent. Therefore the appliance capacity required at each member site is reduced and the bandwidth required for synchronizing data across members is minimized.

   This intelligent partitioning and replication also minimizes the replication load on the master.
4. Provides Grid-wide monitoring and reporting: The Master serves all monitoring and reporting data to the Infoblox Grid Manager application, including the status of member devices, real-time and historical views of DNS and DHCP data, and service logs.

Grid Masters are typically deployed in high-availability (HA) pairs at key data centers, network operations centers, and disaster recovery sites. Any Infoblox appliance (or HA pair) in a Grid can be designated as a Grid Master, as long as it has sufficient database capacity to support all of the unique database objects in the Grid. Database objects include hosts, static IP addresses, dynamic IP addresses, and other network services data elements.

Any appliance (or HA pair) in the Grid, with capacity equal to or greater than the Grid Master, can be designated as a "Master candidate." At any time, an administrator can promote a Master candidate to become the Master. Once an appliance has been newly designated as a Master, it takes just a few minutes for the appliance to contact all member appliances and assume the master role. The flexibility of assigning the Master role and the speed with which a newly designated Master can communicate with affected appliances makes disaster recovery and "follow-the-sun" management transfers extremely natural and easy to do.

Secure Communications and Atomic Transactions

To ensure data privacy and authenticity, all communications among Grid members and the Master are secured with certificate-based authentication and SSL encryption. All data, including DNS zone data, is secured in this manner when being replicated to members of an Infoblox Grid.

Given the critical nature of core network services, it is essential to ensure that they are always available and that the data is always correct and consistent. To ensure the integrity of network services data and prevent data duplication, loss, or corruption, data updates to and from members and the Master—and between members in HA pairs—use atomic database transactions.

An atomic database transaction encompasses a series of tasks that must each execute successfully in order for the transaction to complete. If an error prevents any portion of an atomic transaction from completing successfully, the entire transaction is "backed out" to prevent the databases in different appliances from being left in an inconsistent state.

The banking industry utilizes atomic transactions to maintain the integrity of bank transactions. For example, atomic transactions ensure that a customer's bank balance is always the same no matter which automated teller machine (ATM) is used to access the account. Atomic transactions also protect accounts if a sudden loss of communication or device failure occurs.

In distributed network services systems that do not use atomic transactions, common situations such as failover from an active to a backup device in an HA pair will nearly always result in data inconsistencies. This causes vexing problems, including:

- Issuing identical IP addresses to multiple devices
- A DNS entry pointing to the wrong host (which renders key servers and applications unavailable)

Similar problems can occur if there are failures part-way through a configuration update, such as changing the IP address, subnet mask, gateway, and other IP settings associated with a host. If only some of the configuration parameters are changed and then a failure occurs, the device may receive corrupt configuration data and may become unreachable.

Real-time Data Updates

Users and devices in IP networks are increasingly mobile and transient, resulting in frequent changes in network services data. Prior to mobile computing and wireless access, a typical desktop might have had the same IP address for months at a time. Today, as users move their laptops, IP phones, and other devices from room to room and between different wireless access points, their IP addresses may change several times per day. These changes must be reflected immediately and consistently across a distributed enterprise. This is necessary so that, for example, if a consultant is terminated and their laptop's MAC address is removed from the list of "allowed" devices, the change is reflected *immediately* across all DHCP servers in the enterprise.

In an Infoblox Grid, local changes (such as issuing a DHCP address, renewing a lease, or receiving a Dynamic DNS update) are propagated immediately from members to the Master and vice versa. In addition, the database in the Master is not a delayed snapshot of what was happening in the network at remote sites some time in the past. Rather, the database in a Master (or Master candidate) reflects the real-time state of all of the data across all of the appliances in the Grid at that moment. As a result, status monitoring and data reports, which are served by the Grid Master, always reflect the real-time state of the network.

Ensuring Business Continuity

The Infoblox Grid makes it especially easy and cost effective to support business continuity by maximizing the availability of services and minimizing time-to-recovery in the face of myriad failure scenarios. Example failure scenarios include:

1. Loss of connectivity between a member and the Master: The member device will enter the "disconnected operation" state, in which it will continue to provide all services and will queue updates bound for the Master. When connectivity to the Master is restored, the member will automatically propagate all queued updates to the Master. Once updated data is received by the Master, it will synchronize all appliances in the Grid (including the member that was temporarily disconnected).
2. Failure of an appliance in an HA pair: The backup appliance in an HA pair will detect the failure of the active device within five seconds using industry-standard Virtual Router Redundancy Protocol (VRRP) and will start responding to DNS, DHCP, and TFTP requests within that period. The appliances in the HA pair share a common virtual IP (VIP) address so the transfer

of the passive device to active service is transparent to all clients. Transactional integrity of the updates between the active and backup appliance ensures that the backup appliance's database is always an exact copy of the active device, ensuring, for example, that no duplicate IP addresses are issued by the backup device following a failover.

3. Replacement of a failed appliance: Any like appliance can be used to replace a failed appliance. Here are the steps that occur when a like appliance replaces a failed one, all of which takes place automatically:

    - The like appliance is configured with the IP address of the failed device
    - The like appliance establishes connectivity with the Master.
    - The Master checks the version of software on the replacement member unit.
    - The Master will download and upgrade the appliance software to the version running on the Grid.
    - The Master will load all configuration and service data and will start services running on the replacement appliance.

    This process supports deployment of new appliances at remote sites where there are no skilled personnel.

4. Loss of the Master: If the Master (or Master HA pair) should fail or become unreachable due to a WAN failure or general data center failure, all member appliances will enter the disconnected operation state and will continue to serve data. At any time (either after or before loss of a master) an administrator can contact a master candidate and issue a "promote to master" command. The Master candidate will assume the role of Master and will contact all members informing them of this change.

    If this action is taken before a Mmaster is lost, the Master candidate's database will contain an identical copy of the master's database, so the time required to re-synchronize the master candidate and the members will be minimal.

    If the master promotion takes place after the master fails and the member devices have entered the disconnected operation state, the newly promoted master will automatically re-synchronize the Grid—which can occur in a matter of seconds depending on the total number of objects in the database, the bandwidth of WAN links, and the number of changes that occurred during disconnected operation. However, at no time is service interrupted on the member devices—all synchronization activities are invisible to users. The Infoblox Grid maintains nonstop local service delivery and provides a seamless, fast, automatic way to recover central control and reporting.

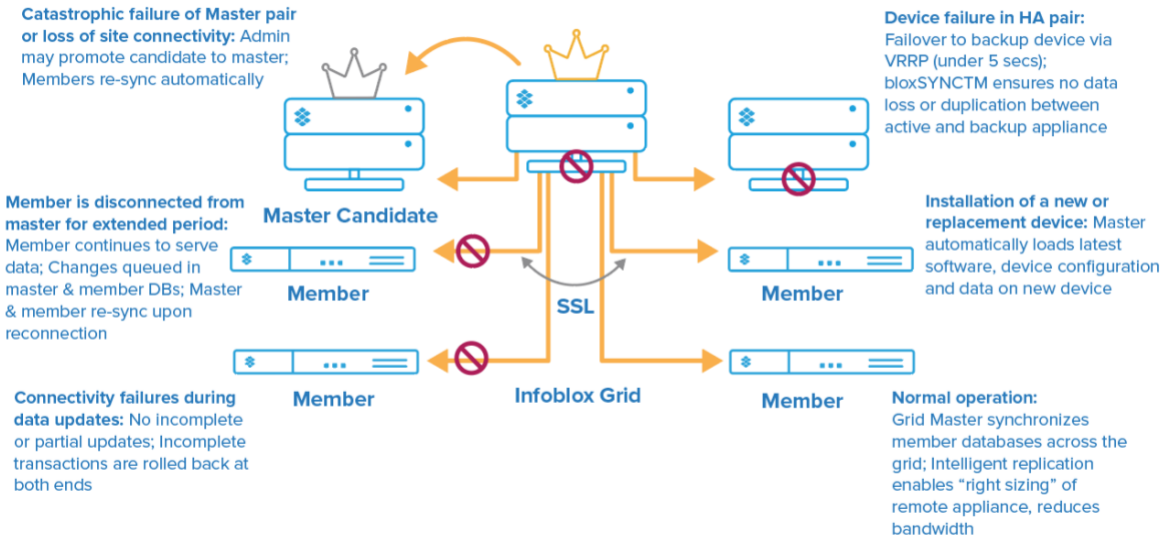

Figure 2: Infoblox Grid Technology assures nonstop services and data integrity in the face of myriad potential failure scenarios

First- and Second-Generation Approaches are Flawed

Legacy solutions for DNS, DNSSEC, DHCP, and IPAM exhibit significant limitations. The most popular legacy approach has been to deploy DNS and DHCP services on "white-box" servers running Linux, Unix, or Microsoft Windows operating systems using software such as ISC BIND or DHCP and Microsoft Windows Server software. This approach is insecure, expensive, and unreliable because it depends on servers with vulnerable, standard operating systems and is susceptible to service disruptions from simple mistakes in manually-edited configuration files.

Deploying, securing, and managing servers at remote locations is expensive and time-consuming. Providing for high availability and disaster recovery and eliminating data loss and corruption is impractical since legacy solutions do not have this functionality built-in. Coordinating activity and correlating data across multiple network services is next to impossible.

First-generation network services—introduced by Infoblox in 1999— addressed security and management challenges by delivering network services in security-hardened appliances with easy-to-use administrative interfaces and highly desirable features such as built-in support for high-availability failover. This approach works well for applications that require a relatively small number of appliances, such as external DNS services, as well as for DNSSEC services delivered within a datacenter. However, first-generation appliances are managed one at a time and do not consolidate network services data to provide enterprise-wide reporting.

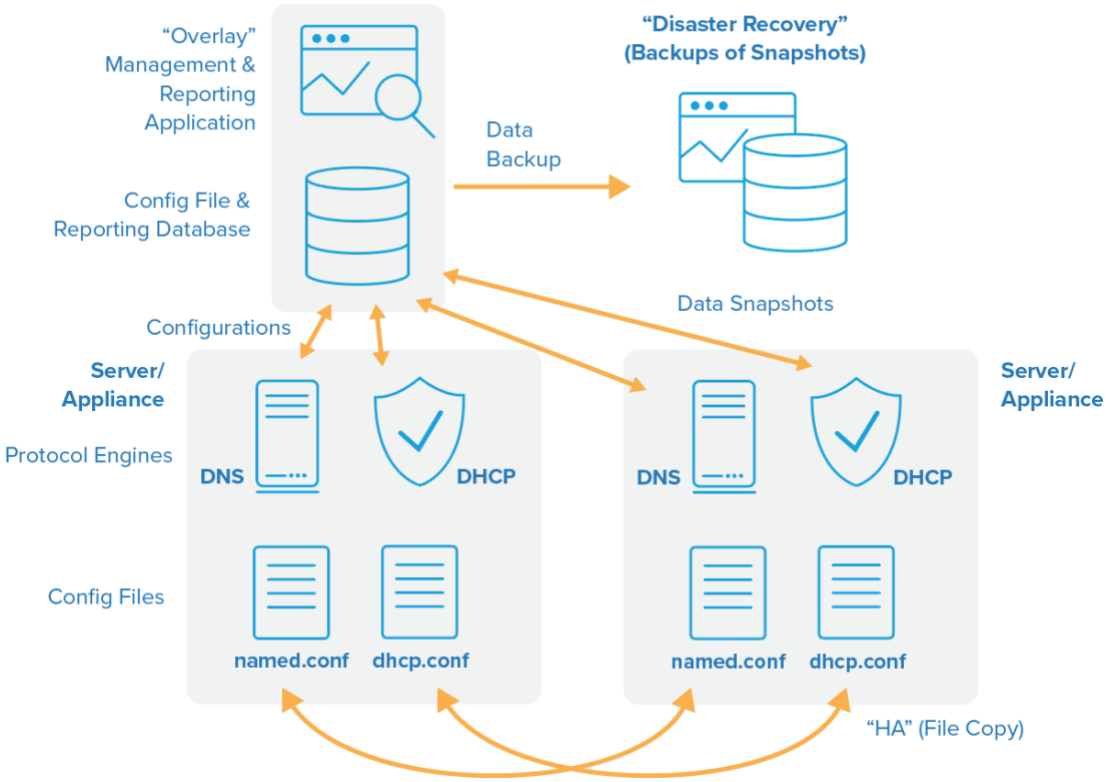

Figure 3: Second-generation "IPAM" systems exhibit significant limitations

Second-generation core network services appliances were introduced to simplify the management of multiple distributed appliances and to provide unified reporting and IPAM capabilities. The second-generation approach mirrors legacy, software-based IPAM systems, which use a separate database and an "overlay" management and reporting application. The overlay application is used to generate configuration files and push them to remote appliances, and is also used to periodically "scrape" data from remote appliances and collect these data snapshots into a dedicated database for reporting.

However, there are several problems with this approach. Most importantly, file- based systems lack transactional integrity, and so they are susceptible to data loss and corruption. For example, file-based systems for DHCP often lose synchronization during routine high-availability failovers, resulting in IP address conflicts and naming inconsistencies that can be extremely difficult to troubleshoot and correct.

In addition, IPAM information and reports from second-generation systems can only approximate the data actually being served in the network at any moment, because the dedicated IPAM database is assembled from periodic snapshots of the data on the remote appliances. This time-delayed approach may have been acceptable for yesterday's static networks; however, today's dynamic networks, with applications such as wireless and dynamic addressing, exhibit frequent changes in DNS and DHCP that must be reflected in real-time.

Finally, disaster recovery for second-generation systems is a "built-on" rather than a "built-in" function, typically requiring manual management of data replication to a backup database, and may even require manual "re-homing" of remote servers to point to the disaster recover site in the event of a failure.

Today's requirements for network security, availability, and compliance demand that naming services, address assignments, network access decisions and usage reports be based on an accurate, authoritative, real-time view of which users, devices and addresses are in use. For this reason, second-generation systems are inadequate.

Infoblox Grid: A New Generation of Technology

Unlike first- and second-generation solutions, Infoblox Grid technology assures nonstop availability of distributed core network services with full data integrity and real-time reporting. Infoblox Grids are implemented by securely networking together the databases embedded within each appliance. The database integrates and correlates a wide range of network services data elements, including IP addresses, host names, devices addresses, and even firmware images and configuration files for IP phones and other devices.

Changes to the data that occur on any appliance are reflected across the Grid, securely, in real time and with full transactional integrity. This prevents data loss, eliminates possible inconsistencies and errors and ensures that usage reports, address assignments, and network access decisions are based on accurate data. Because the Grid does not require a separate, external database for device configurations and reporting data, Infoblox Grids provide inherent reliability advantages, data integrity, faster and easier disaster recovery, and are easier to scale and manage compared with legacy or second-generation appliance approaches. Figure 4, below, demonstrates the flexibility of Infoblox Grid.

Infoblox Grid Manager

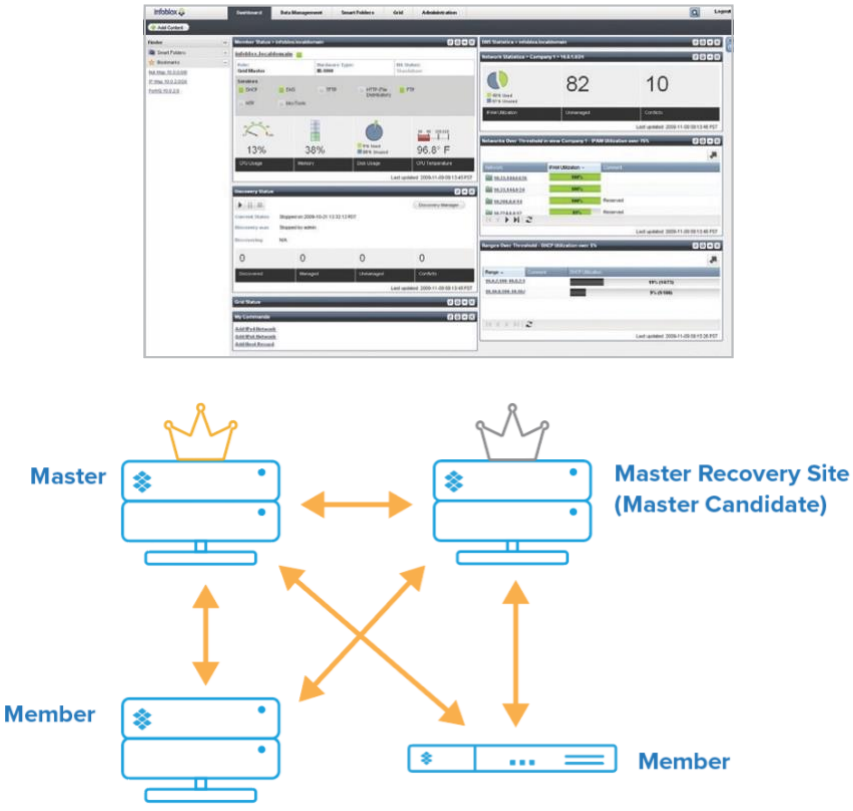

Figure 4: Infoblox Grids are networks of distributed databases

Infoblox Grids solve immediate needs for distributed enterprises, and have far-reaching implications as a critical new type of IT infrastructure. For example, the sophisticated software download and remote system monitoring and control capabilities of Infoblox Grids are used to enable appliance-based deployment of remote DNS and DHCP services for the Alcatel-Lucent VitalQIP® IPAM solution. The unique ability to turn TFTP into a centrally managed, distributed service for VoIP deployments is another example of the power and flexibility of Grids. Infoblox Grids represent the next generation of appliance-based core network services delivery, and are an exciting new kind of IT infrastructure for providing nonstop, centrally managed services in distributed environments.

Corporate Headquarters | 3111 Coronado Dr. | Santa Clara, CA | 95054
+1.408.986.4000 | 1.866.463.6256 (toll-free, U.S. and Canada) | info@infoblox.com | www.infoblox.com