NetApp Kilo Client: A Technical Case Study

Product and Partner Engineering | Network Appliance, Inc.
January 2007 | TR-3566

Abstract

This technical case study is based on interviews and hands-on work conducted by Network Appliance SAN Product and Partner Engineering in cooperation with NetApp RTP Engineering Support in Research Triangle Park, North Carolina. The case study examines the NetApp SAN infrastructure deployed for the Kilo Client compute grid.

Table of Contents

1 Introduction 3
2 Background 3
3 Project Objectives 3
4 Key Requirements 4
4.1 Rapid Provisioning 4
4.2 Simplified Management and Configuration 6
4.3 Scalable and Flexible 7
5 Architecture 7
6 Kilo Client Usage Scenario 9
7 Sizing Recommendations 11
8 Summary 11

1 Introduction

The Network Appliance™ Kilo Client is a compute grid consisting of over 1500 compute blades, Ethernet and Fibre Channel networks, as well as NetApp storage systems. It was designed and implemented by the NetApp Engineering Support team in Research Triangle Park, NC. This group supports product development and quality assurance engineering teams across five global NetApp design centers (Sunnyvale, CA; RTP, NC; Pittsburgh, PA; Waltham, MA; Bangalore, India).

This case study details the implementation of the NetApp Kilo Client and how it leverages NetApp iSCSI and Fibre Channel SAN storage technology to simplify the management of a large compute grid. Specifically, the use of NetApp LUN clones and SAN boot are core technologies of the Kilo Client.

2 Background

As compute and storage requirements grow, NetApp has scaled its product line and in 2006 announced high-performance Data ONTAP® GX storage systems as well as performance and scalability enhancements to Data ONTAP 7G. Increasingly the flexibility and scalability of NetApp storage systems are a good fit in grid computing environments. Large technical computing environments often have demanding storage requirements. These environments also have many compute nodes or clients attached to NetApp storage systems. In order for NetApp engineers to test scalable and flexible NetApp environments, a large compute grid environment was necessary. The Kilo Client project was initiated to address these needs.

3 Project Objectives

The key drivers for the Kilo Client project were NetApp Engineering needed ways to test increasing product limits, generate more rigorous load and stress tests to determine the breaking points, and develop and prove NetApp scalability and flexibility for grid storage provisioning technology.

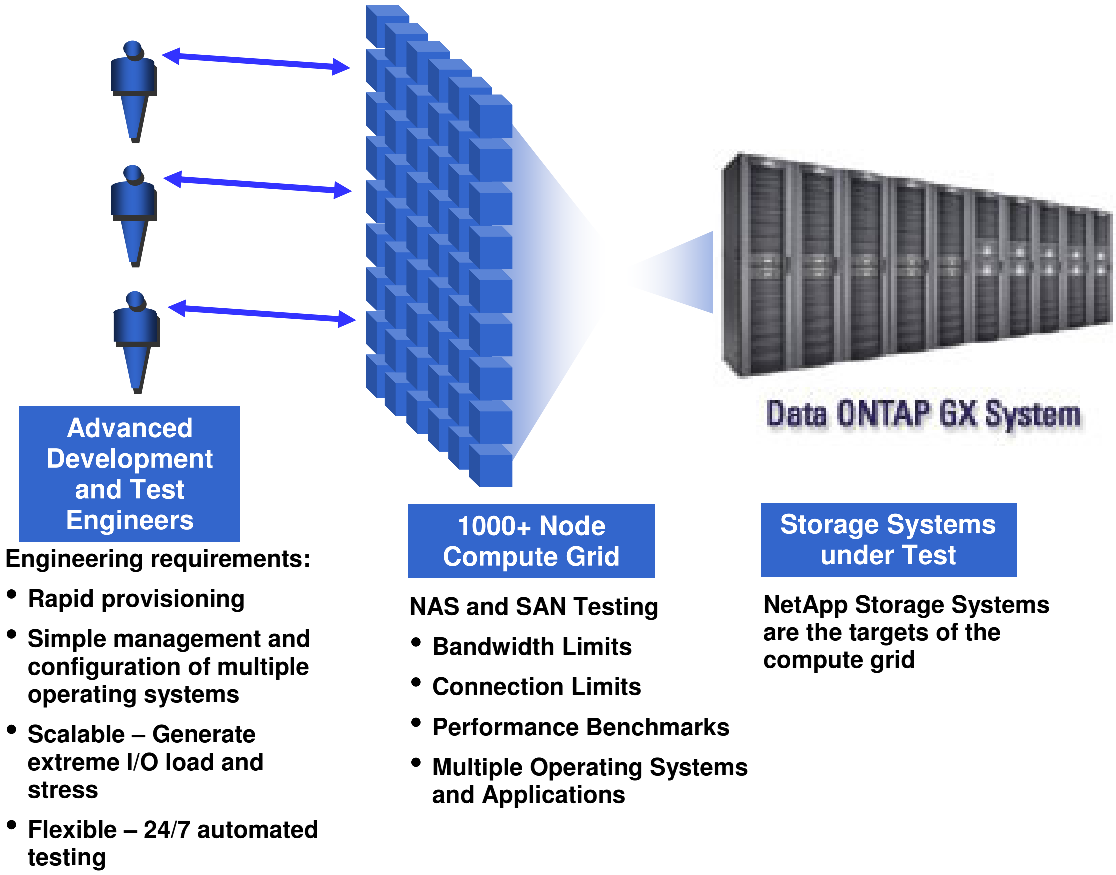

Figure 1) NetApp Kilo Client overview.

4 Key Requirements

The Engineering Support Team in conjunction with engineering development and test teams determined that for the Kilo Client to be a useful engineering tool it had to meet the following requirements:

- Rapid provisioning
- Easy management and configuration
- Scalable and flexible

4.1 Rapid Provisioning

With over one thousand nodes in the Kilo Client, every minute it takes to provision an operating system counts. Full disk OS provisioning tools were evaluated, but it took many minutes to install an OS image on a single node, which turned into hours for a large compute grid. The full disk OS provisioning option also added the management complexity of dealing with boot image servers and scaling the servers to handle the bandwidth requirements for imaging hundreds of nodes simultaneously. Using SAN boot capabilities from a NetApp FAS storage system along with LUN clone and FlexClone™ technology allowed the engineering support team to rapidly clone OS boot LUNs for all the Kilo Client nodes.

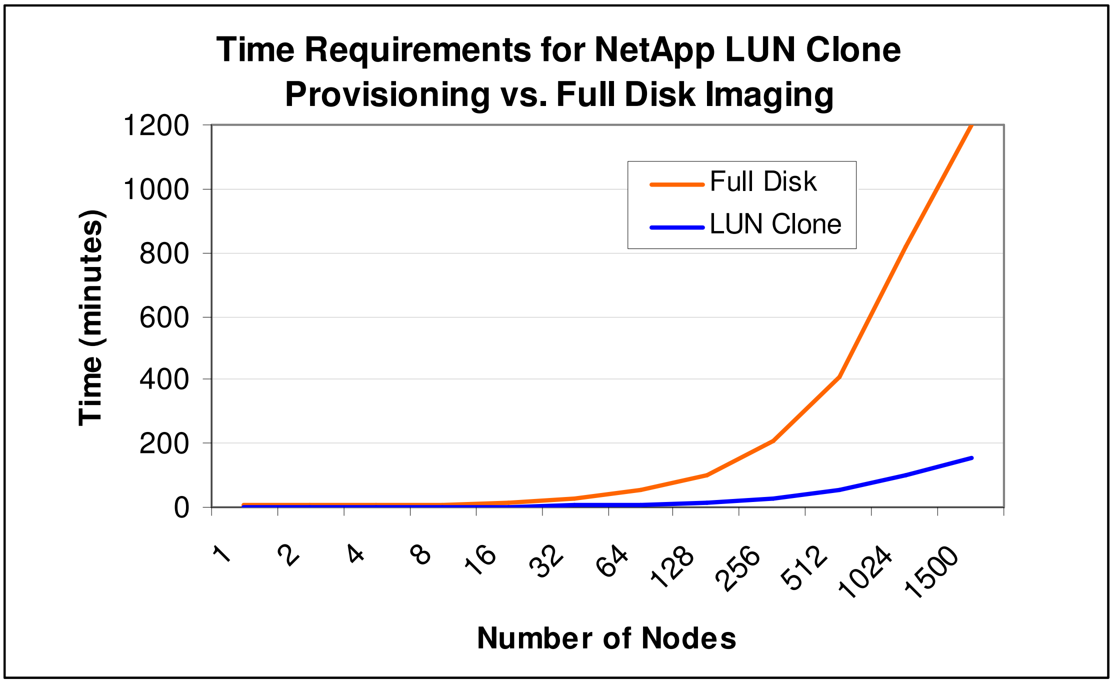

Figure 2) Time required to provision Kilo Client nodes with LUN clones versus full disk imaging.

Additionally, the LUN clone technology used to rapidly deploy OS boot images has the benefit of increased disk space savings through flexible provisioning. With NetApp LUN clones only the data that has changed on the OS LUN consumes additional physical disk space. Typically boot disks have static data with the only difference between boot images being the node-specific configuration information (network settings, unique system identifiers, and so on). The NetApp LUN clones are backed by the original OS image LUN and therefore only consume a fraction of additional storage space with the specific configuration data.

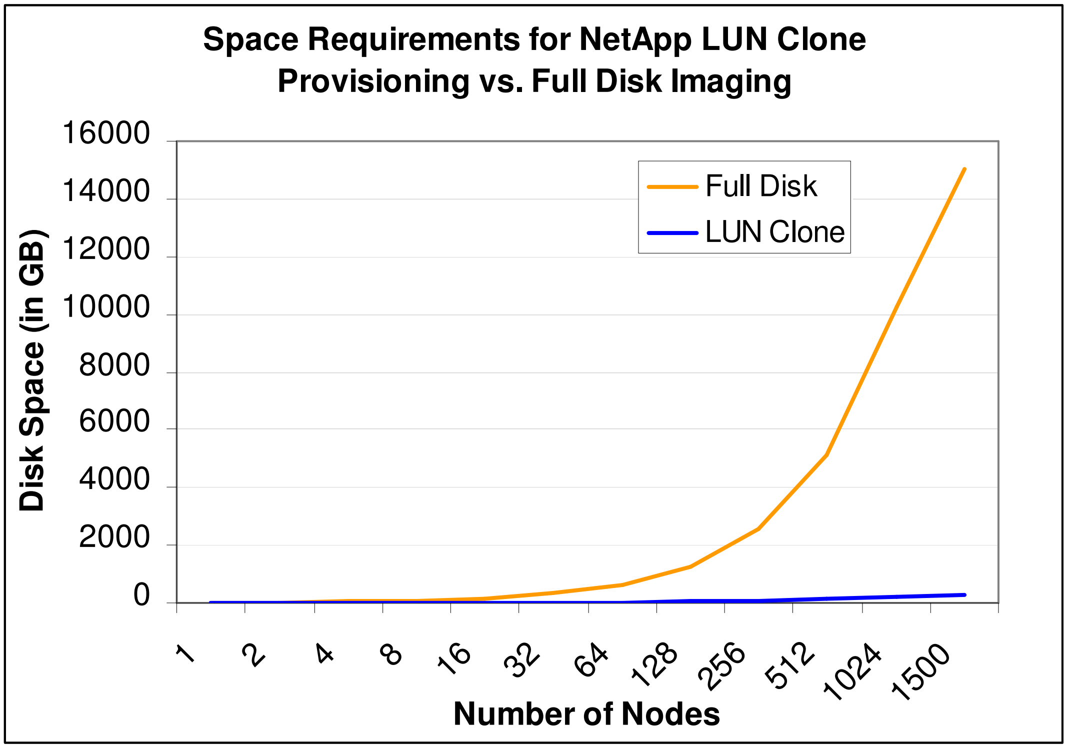

Figure 3) Disk space required to provision Kilo Client nodes with LUN clones versus full disk images.

4.2 Simplified Management and Configuration

NetApp QA engineers require numerous nodes to generate load and stress on NetApp FAS systems under test. The host OS configurations vary depending on the types of tests being run. Managing each host OS image with patches and applications independently was unwieldy. The solution was to create a “Golden LUN” OS image with the specific engineering requirements, then use NetApp LUN clone technology to present flexible provisioned writable LUN copies to each Kilo Client node. Moreover, the nodes in the Kilo Client only have to be configured once to SAN boot from the NetApp FAS system. The OS boot LUNs are managed from Data ONTAP and can be easily mapped and unmapped according to which OS needs to be booted on a node.

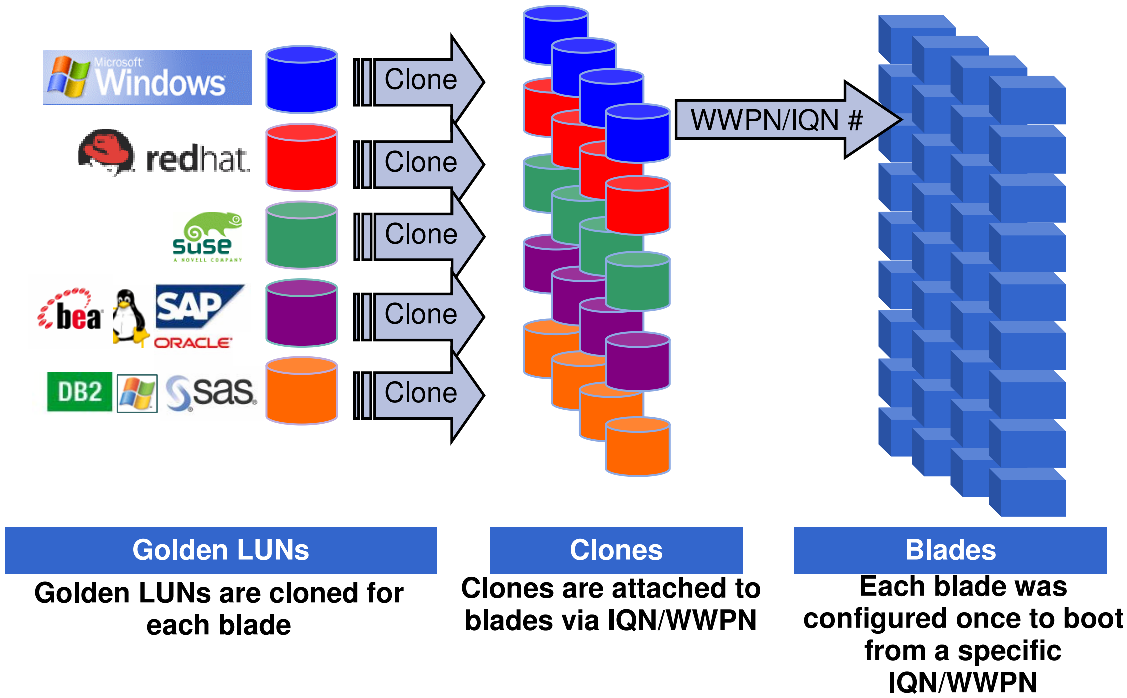

Figure 4) NetApp Kilo Client golden LUN configuration.

4.3 Scalable and Flexible

The initial Kilo Client installation of 1000 nodes was just the beginning. There were no plans to stop at just 1000 nodes. In just over half a year the Kilo Client has grown 50% to 1500 nodes to date. This rapid deployment is due in large part to the NetApp SAN storage infrastructure it relies on. The scalability and flexibility of the NetApp FAS systems used to serve the Kilo Client boot images allow for additional nodes to be brought online within minutes of racking the hardware. The Kilo Client nodes can also be easily split into small subsets of nodes for use by multiple engineering teams or different combinations of operating systems, and applications may be allocated to a particular test.

5 Architecture

The NetApp Kilo Client architecture consists of blade nodes, network/fabric infrastructure, NetApp FAS systems for boot LUNs, NetApp NearStore® systems for boot LUN archival, and the NetApp FAS farms under test.

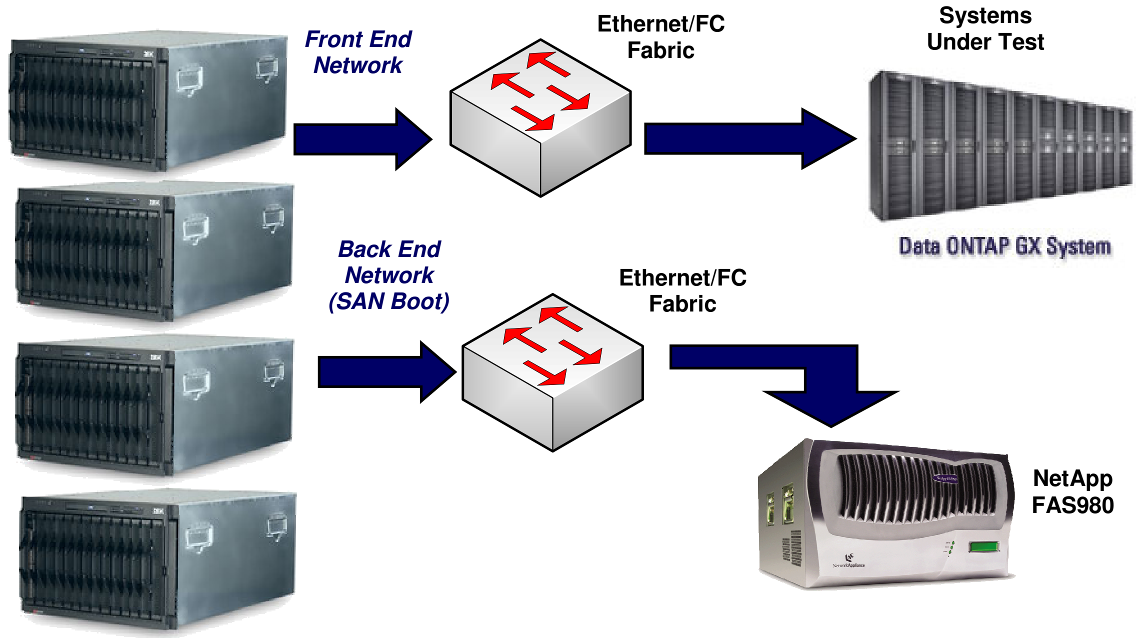

Figure 5) High-level NetApp Kilo Client architecture.

The Kilo Client nodes are contained in IBM BladeCenter chassis with the HS20 Intel® blades as the current majority of the blades. Some POWER JS20 and AMD LS20 blades are also part of the Kilo Client infrastructure.

The network infrastructure consists of three separate networks – one for iSCSI boot, one as a management interface to the blade clients, and the network dedicated to the NetApp systems under test. The iSCSI boot network is a Gigabit Ethernet network dedicated for iSCSI traffic to the boot LUNs on a NetApp FAS980 storage system.

Figure 6) NetApp Kilo Client podule.

6 Kilo Client Usage Scenario

The RTP Engineering Support team is responsible for preparing and managing the Kilo Client for use by NetApp engineering development and test teams. The RTP Engineering Kilo Client support team created a set of base operating system “golden LUN” images during the initial roll-out of the project. The golden LUN images currently include versions and variants of Windows® Server and Linux® operating systems. These “golden LUN” images reside on the FAS980 storage system and are only booted when the image needs to be updated. For example, when new Microsoft® updates are released for Windows, the golden LUN is booted, and the updates are applied. As an extra level of protection, the golden LUN image resides in a Snapshot™ copy and has been mirrored to a NetApp NearStore system for quick retrieval in case the update corrupts the golden LUN image.

When a NetApp engineer requests time on the Kilo Client, the engineering support team first provides the engineer with one blade system SAN booted from a LUN clone of the golden LUN. The engineer then customizes the SAN booted operating system by installing applications and test tools. Once the engineer has completed the customization of the single blade boot LUN, a Snapshot copy of the LUN is created. The Snapshot copy of the boot LUN can then be mirrored to the other NetApp FAS storage systems in the Kilo Client that provide SAN boot services. The original customized boot LUN is then cloned using the Data ONTAP LUN clone technology. The LUN clone is split from the original LUN, and it is unmapped from any blades to preserve the original LUN and may be archived to a NearStore system for future use. Now a Snapshot copy of the FlexVol® volume containing the OS SAN boot LUN is created, and writable LUN clones are created for as many blades as the engineer wants to use for testing.

The LUN clone process has the immediate advantages that it takes seconds to create the LUN clone, and it's thin provisioned. Thin or flexible provisioning of the boot LUN in this environment allows each bootable LUN clone to be unique and writable yet share blocks from the original LUN. Several hundred LUN clones may be created from the original LUN as the OS boot disks for the blades in the Kilo Client. It takes minutes instead of hours for the Kilo Client Engineering Support team to provision hundreds of Kilo Client systems for engineering test. And finally, when the engineer's reservation for the Kilo Client blades expires, the engineer may request that the configured boot LUN be kept on a NearStore system to save time configuring the test client the next time the engineer has a reservation.

7 Sizing Recommendations

There are several important best practice guidelines that NetApp currently recommends for compute grid SAN boot environments:

- Use of a local disk for OS paging/swap space is highly recommended for all FC, iSCSI HBA, and software iSCSI SAN boot environments. It is required for iSCSI software SAN boot environments.
- Set the FlexShare™ priority of volumes with SAN boot LUNs to “high” or “veryhigh.”
- Use the NetApp Fibre Channel and iSCSI Configuration Guide for supported topology recommendations and general sizing guidelines.
- SAN boot system sizing maximums:

Controller	FAS270	FAS3020	FAS3040	FAS3050	FAS3070	FAS6030	FAS6070
Number of SAN Boot Hosts	16	128	128	128	252	252	252

8 Summary

The Network Appliance Kilo Client has successfully proved itself as a valuable tool for NetApp engineering teams as well as an example of leveraging NetApp SAN innovation to address the challenges of grid computing.