



# DEPLOYING INNOVATIVE MESSAGING SERVICES USING OPTIMIZED SYSTEM ARCHITECTURE

An Infinite Convergence and Dot Hill White Paper

April 2014



## Message Storage Application (MSA) Solution Brief:

*Infinite Convergence deploys innovative Telecommunications Messaging Products leveraging flexible Storage Solutions designed for scalability.*

*Deploying cost-effective, innovative messaging products that scale to customer needs while maintaining five 9s availability demanded that Infinite Convergence find a storage solution that was optimized for scalability, flexibility and price. Partnering with Dot Hill and using Cassandra provided Infinite Convergence with a solution that meets the high expectations of the Telecommunication marketplace.*

### Overview

Operators are faced with declining messaging revenue as subscribers turn to Over-The-Top (OTT) services as their choice for exchanging mobile messages. According to Ovum, OTT messaging apps will have cost operators \$32.6 billion in lost SMS revenue in 2013, rising to \$86 billion in 2020. The Message Storage Application (MSA) from Infinite Convergence Solutions provides operators with new innovative messaging services that recapture the attention of the subscriber base.

The MSA allows subscribers to start messaging on one device, like a mobile phone, and continue the conversation on another device, like a laptop computer or tablet, using one mobile number. There's no need to remember multiple numbers for the same user, regardless of which device they may be using. All messages are seamlessly synchronized in real-time across all devices, allowing the subscriber to use his preferred devices for exchanging messages. Message content is stored in the network for easy access. The MSA supports feature-rich capabilities beyond text and picture messaging, including group messaging, location sharing and searching.




Figure 1: Cellular Network that includes the MSA



## Challenges:

Relying only on a rich feature-set is not sufficient to launch a successful messaging service. A successful service must provide a combination of innovative features along with the core attributes of reliability, high availability, scalability, and high performance.

In January 2013, a large carrier<sup>1</sup> announced the delay of its RCS service due to “complex integration issues” and “lack of stability.” It may be tempting to focus first on feature functionality. But if the system crashes, loses messages, or is hard to provision, then the service is bound to fail. It is important that the vendor for the messaging service has a proven track record of delivering and deploying stable systems that can scale to accommodate uptake by subscribers.

The days of waiting for several minutes to send or receive small pictures and videos are over. Digital cameras on phones generate high-resolution pictures and videos and consumers want to share them in their original size. They want their friends to start viewing their videos the minute they press the send button. A low latency user experience and high performance are essential. Any messaging service should be able to handle extreme load both in terms of high message rate and size of messages transferred while maintaining low latency.

Infinite Convergence wanted to efficiently manage a range of subscribers from 0.5M to >100 Million while synchronizing messages across multiple devices in under 200 milliseconds (ms)! The goal was to have the subscriber put down his mobile phone and pick up his tablet and instantly have all the message history already there regardless of OS platform. A solution that can scale to over 1,000,000 TPS (transactions per second) with less than 200 ms latency and over 99.999% availability demands big data techniques.

To achieve quick access to the message database, as well as enable search, message expiration, multi-user and group access to messages, Infinite Convergence designed a proprietary message metadata database. The metadata database required using a “NoSQL” style database that needed to be maintained on large arrays of disk storage – designed for speed, data protection, reliability, availability, scalability and manageability.

Infinite Convergence chose Cassandra to store the message metadata content. This database of ‘metadata’ is used to store indices of the messages in the long-term data store, as well as associated user information. The MSA follows common usage patterns expected by Cassandra, and thus can be deployed using standard practices: commodity hardware, massively parallel processing, and shared nothing.

The metadata storage access operation requires numerous Random Read operations, which are expensive and incur additional load on the storage subsystem. To compensate, the solution needed to deploy faster storage for the metadata. Some of the options that were considered:

---

<sup>1</sup> [http://www.telecomtv.com/comspace\\_newsDetail.aspx?n=49947&id=e9381817-0593-417a-8639-c4c53e2a2a10](http://www.telecomtv.com/comspace_newsDetail.aspx?n=49947&id=e9381817-0593-417a-8639-c4c53e2a2a10)


<sup>2</sup> 5 9s (99.999) availability means a system downtime of less than 5.39 minutes annually.

- 1) Multiple high speed (10K RPM or 15K RPM) hard disk drives (HDD), striped together using RAID 0 or RAID 10 (providing greater aggregate Random I/O performance)
- 2) A solid state drive (SSD) or Flash card instead of HDD for storage. SSDs provide as much as 100x the Random I/O performance of a traditional HDD.


The shortcoming of the first option above was the sheer quantity of drives needed to deploy and manage each node simply to get performance to an acceptable level. The shortcoming of using SSD or Flash was the expense of deploying them on every node. To overcome these issues, Dot Hill introduces an alternative design that addresses both concerns.

### Storage Solutions for Messaging Services

Dot Hill is introducing an alternative to hanging more expensive storage on every physical server node in a Cassandra Cluster to address the performance shortcomings of running a metadata database in a messaging service. The solution uses a Dot Hill AssuredSAN Pro 5000 storage system for each 'Rack' group of nodes within a cluster. The 'Rack' group would share the storage from a single AssuredSAN Pro 5000. Cassandra redundancy policies are maintained by defining the 'Rack' within Cassandra, and ensuring all replications of node data occur between nodes that are not of the same rack. See figure 2 for an example of a 'Rack' configuration.



**Figure 2**  
Mobile Messaging Metadata Store



In the 'Rack', the Cassandra compute nodes consist of commodity servers, but include iSCSI HBAs to connect to the shared storage. A pair of Ethernet switches is necessary to provide redundant connectivity between the storage and the servers. The AssuredSAN Pro storage array will be configured into two pools of storage. The disks are grouped together in RAID 1 or RAID 10 arrays to provide the optimum redundancy and random I/O performance.

For each node, one or more volumes are allocated from the storage pools in the array and presented to the node. Data that is stored on the volumes is automatically migrated between SSD and HDD in real time by the array controllers. Balancing the storage location for the metadata can dramatically improve overall system performance.

## Solution Benefits

Why is the AssuredSAN Pro 5000 a preferred solution for a Messaging Metadata Store? The primary benefit comes from the performance and value advantages offered by the array. The AssuredSAN Pro 5000 is a multi-tier storage device that includes sophisticated algorithms to automate the placement of data in each tier to optimize performance. The highest tier of storage is SSD. Frequently accessed blocks of data are stored here, providing extremely fast read access. The lower tier of this solution is 10K RPM SAS drives, providing capacity and value for the data that is less frequently accessed. The overall solution provides the perfect balance between high performance hot storage, and inexpensive cold storage.

Other benefits include:

- 1) Real-Time Storage Migration. The algorithms to identify hot data blocks and migrates them to the SSD tier operate in Real Time. This is ideal for Messages and Message metadata that are relevant for a short period of time.
- 2) Sequential workloads stay on lower tiers of storage, improving performance. In addition, the system is designed to handle many sequential operations simultaneously. These features are ideal for handling the intensive sequential write operations from many Cassandra nodes at the same time.
- 3) Volume Thin Provisioning. No need to guess which node needs the most capacity. Each volume takes only what it needs from the array.
- 4) Scale up to size for capacity. JBOD expansion chassis provide more capacity as required by the solution. Storage can be added dynamically without impacting the operation of the compute nodes.
- 5) High availability (HA) features. While Cassandra is designed to allow for node failure or maintenance activities, the inclusion of RAID protection allows for simplified management of the components that are most prone to failure.

## Final Remarks

The Infinite Convergence MSA enables centralized storage of messages in the network and synchronizes them across multiple subscriber devices. The design of the software and the overall MSA system architecture, including its scaling and its reliability functions that are engineered into the architecture, is



critical to achieving the type of experience that truly permits users to move freely from one device to another and maintain context while devices are instantaneously synchronized. The Dot Hill AssuredSAN Pro 5000 is a key component of that successful product solution. Currently, Infinite Convergence has deployed the MSA in a North America Tier-1 Cellular environment, proving the well-designed software and hardware architecture that results in successful large-scale commercial deployments.

The Dot Hill AssuredSAN Pro 5000 offers a unique combination of features and advantages that make it an ideal candidate for storage subsystems in carrier-grade Telecommunications Messaging services. With a combination of high-performance SSD drives, low- cost HDD drives and real-time migration of blocks to the proper tier, the Pro 5000 blends state-of-the-art technologies together to deliver performance, availability, protection and value to storage solutions.

## About Infinite Convergence Solutions

Infinite Convergence Solutions is a provider of innovative and cutting edge products supporting the needs of next-generation wireless communication. Our Messaging Business provides a comprehensive Messaging Product Portfolio that includes Rich Communication Suite (RCS), Message Storage Application (MSA), Converged Messaging, MMS Center, SMS Center, SMS Gateway and Enterprise Messaging Service. Our messaging solutions are scalable, highly available and offer industry leading performance. They are deployed in Tier 1 and Tier 2 mobile operators as well as Enterprises world-wide.

## About Dot Hill

Leveraging its proprietary Assured family of storage solutions, Dot Hill solves many of today's most challenging storage problems – helping IT to improve performance, increase availability, simplify operations, and reduce costs. Dot Hill's solutions combine breakthrough software with the industry's most flexible and extensive hardware platform and automated management to deliver best-in-class solutions. Headquartered in Longmont, Colo., Dot Hill has offices and/or representatives in China, Germany, India, Japan, Singapore, the United Kingdom, and the United States.

## Contact Us

### Infinite Convergence Solutions

Website [www.infinite-convergence.com](http://www.infinite-convergence.com)  
Twitter @nxtgenmsg  
YouTube Infinite Convergence

### Dot Hill

Website [www.dothill.com](http://www.dothill.com)  
Twitter @Dot\_Hill  
YouTube DotHillSystems

© 2014 Copyright Dot Hill Systems Corporation and Infinite Convergence Solutions, Inc. All rights reserved. Dot Hill, the Dot Hill logo, and AssuredSAN are trademarks or registered trademarks of Dot Hill Systems. Infinite Convergence Solution and the Infinite Convergence logo are trademarks or registered trademarks of Infinite Convergence Solutions. All other trademarks are the property of their respective companies in the United States and/or other countries.

