



## Powered by the Scality RING, Comcast Serves Over 30 Million Mailboxes with Greater Flexibility and Resilience

Comcast Cable is one of the world's largest video, high-speed Internet and phone providers to residential customers and businesses. Comcast has invested in technology to build an advanced network that delivers among the fastest broadband speeds and brings customers personalized video, communications and home management offerings. Comcast is also one of the largest email environments in the world, with more than 30 million mailboxes.

Comcast was seeking to add additional email storage capacity. At the same time, the company was moving forward with a multi-phase transformation toward a more cloud-driven, stateless, and efficient infrastructure. Scality was able to help support an important component of that work.

An email storage environment of this size has unique characteristics:

**Massive Scale** – billions of objects and multiple petabytes

**Relentless Workload** – high-churn small and large files

**Complex Application and Data Relationship** – geographically distributed users, data availability and durability requirements

Prior to 2013, Comcast was running email server application Zimbra 6 and storing email using a block storage SAN (storage area network). Comcast used hierarchical storage management (HSM) to migrate and manage the data across a complex arrangement of five tiers of storage.

*In order to remain knowledgeable of best practices in storage technologies, Comcast had been testing a number of new solutions, including running a small production email system with the Scality RING since 2011.*

« True elastic scaling becomes attainable allowing for the addition and removal of resources, while the system is running, even during busy hours. »

## Rethinking Email Storage for Petabyte Scale

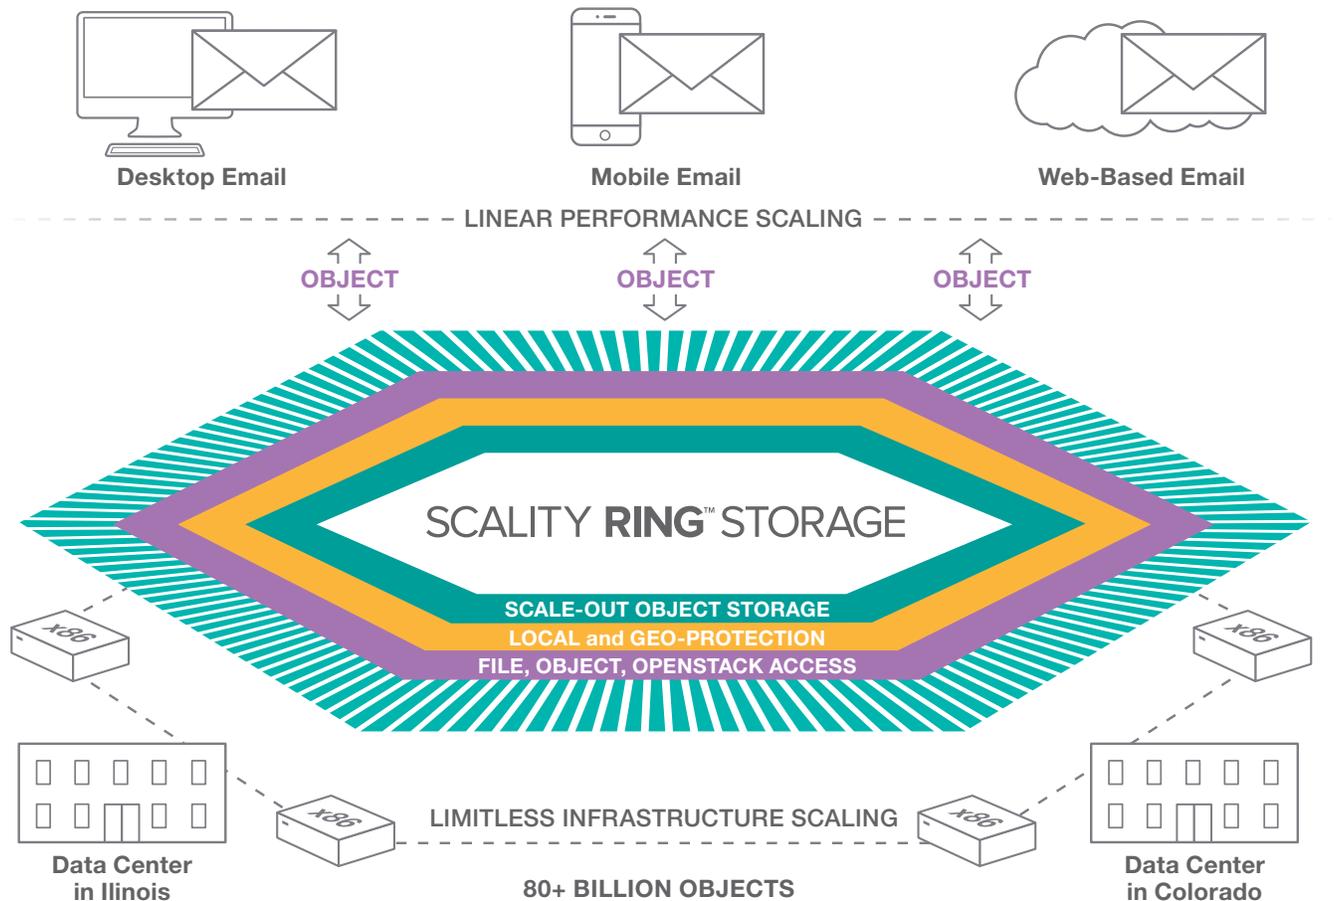
As part of its normal growth and hardware refresh process, Comcast took the opportunity to rethink and reengineer the data center infrastructure powering its email platform. Confronted with petabyte-scale, simplifying storage management processes, and other goals, Comcast targeted a shift to object storage for Binary Large Objects (also known as BLOBs – the messages and attachments themselves).

After careful consideration of its operational needs and storage requirements, Comcast developed a robust set of goals it sought to achieve with a new storage solution, including:

- > Efficient cost-per-gigabyte, to support growing data volumes
- > Ample storage capacity, massive scalability, and easy expansion
- > Performance, resiliency and durability improvements
- > Independence between BLOBs and mailbox metadata
- > More efficient operational processes and reduced effort to operate the platform

## Implementing the Scality RING

Comcast's move to a new storage solution involved consultation and experimentation with the Scality team to strike the best balance between all of Comcast's goals. For example, an early deployment was done on smaller servers with 10x3TB data drives, with a later expansion on higher density servers from a different server manufacturer, with 15x4TB data drives. Ultimately Comcast deployed hundreds of data servers with a mix of 3TB and 4TB SATA drives in multiple data centers paired with Scality's connector for Zimbra 8.



## Comcast Achieved its Economic, Operational, and Architectural Goals with Scality

### Hardware Flexibility, Capacity and Rackspace

- > Scality allows businesses to have data storage on standard x86 servers, instead of closed SAN appliances. In addition, Scality is hardware-agnostic, meaning that any hardware can be utilized and/or swapped in and out of the data center without interrupting operations. This flexibility was of great benefit when Comcast switched hardware after initial deployment. Looking to the future, when Comcast continues to expand, the hardware-agnostic design of the Scality RING will allow it to add heterogeneous hardware. With the RING, Comcast can use a wide selection of performance and capacity-dense servers, from a variety of manufacturers, even multiple generations simultaneously.

### Operations

- > Prior to implementation of RING, Comcast regularly moved mailboxes from one SAN-enabled host to another in order to manage capacity, file system directory sizes, file system inodes and available IO. These were complex operations, made more complex as a result of the large number of HSM tiers.
- > With the implementation of Scality RING, Comcast was able to take advantage of a single tier of storage, utilizing low cost LFF near-line drives, which provides both capacity and all the IO required for message storage.
- > Migrating accounts between servers is now a vastly simplified operation of just moving a small amount of metadata with the actual mail BLOBs remaining in place. Moving just the metadata significantly reduces operational impact because the operation takes much less time to complete.
- > The RING has also increased the resiliency of the mail system: there are no single points of failure with RING. In the event that disk failures occur, there is more data distributed across the system to rebuild from. Plus, recovering from a failure and returning to normal operating state is much faster because there are more disks involved in rebuilding the data than in a traditional RAID group.
- > Scality's approach matched Comcast's overall architectural goals to move to cloud-first designs where compute, storage, and IO capacity exist in shared pools, and can be used to address user application or data load dynamically, either in steady state or business continuity situations. With the RING, storage capacity expansion is simply a matter of adding hardware. Comcast can take servers in and out of maintenance, without impacting users, to attend to a variety of maintenance or emergency issues; software upgrades, OS upgrades, etc. These operations then become non-disruptive to the end user, and further help to improve operations as they allow the data center to scale a lot more easily. True elastic scaling becomes attainable allowing for the addition and removal of resources, while the system is running, even during busy hours.

### Advanced Data Protection

- > The RING's erasure coding (EC) technology, provides data protection and resilience. The RING uses policy to determine when to apply EC based on object size – this provides a mechanism to balance IO and capacity needs.
- > Even though email messages are usually small, Comcast ended up with significant capacity savings using EC. Comcast uses EC (14,4), which has only 29% storage overhead, compared to the 200% overhead it would get with replication with three copies.

« *The Scality RING solution helped achieve continued reductions in price-per-gigabyte and price-per-IOPS costs.* »

### Self-Healing

- > Self-Healing – the ability of the system to bounce back from issues, or to reallocate resources to avoid downtime when particular components of the solution fail is an essential design tenet for Comcast’s platforms and services. Scality RING keyspace allocation is server, chassis, rack and data center aware. If a disk on a server fails, the data is rebuilt on the other disks of the same server seamlessly. If a server experiences an issue, the data is still present and other servers will take the responsibility of temporarily hosting new replicas that should have been created on the failed server. After the server is repaired, those new replicas are balanced off to the recovered server. Even in case of an issue that affects an entire rack, access to the data for reads and writes should not be affected due to rack-awareness and the ability of the remaining servers to redistribute the load between them.
- > As a best practice, Scality ensures that the disks are never filled up passed a certain percentage so that the RING will always be able to handle multiple disk-related issues simultaneously and still be able to recover and rebuild the correct number of copies of all the data on the correct servers.

### Dedicated Care Services

- > Comcast enjoys access to Scality expertise throughout the project lifecycle – from design through operational support and maintenance.

### Maintaining a Balanced System and Creating a Stateless System

- > When Comcast added capacity after its initial deployment, creating a heterogeneous hardware environment, it wanted to rebalance mailboxes and data based on server capacity and capabilities. After implementing the new solution with the Scality RING, Comcast is now better able to move mailboxes to balance out the system and infrastructure.
- > Moving forward, Comcast is working to further automate processes so that mailboxes automatically balance themselves in near real-time. In its previous SAN environment, a significant amount of data had to be moved to achieve this, which increased the chances for errors to occur; it’s more susceptible to corruption, fragmentation, or failure. Comcast has found that with the Scality RING these operations are significantly easier and involve less risk, and mailboxes can be moved and balanced with little-to-no data movement.

Since going into production in 2013, Scality’s RING storage solution has been a valuable addition to Comcast’s continued evolution towards a cloud-scale, software-powered infrastructure. Using Scality technology, Comcast has made its email system more flexible, resilient, and high performance, creating a better product for customers, while also realizing operational efficiencies and cost savings.



San Francisco • Paris • Washington, DC • Boston • Tokyo • Singapore • London

[www.scality.com](http://www.scality.com)  [@scality](https://twitter.com/scality)