



**Driving Innovation  
Through the Information  
Infrastructure**

**SPRING 2011**

**Cheating Death In the Datacenter:  
*Leveraging Cloud to Extend the Life of  
Primary Storage Systems***

Presented By:  
David A Jones  
IT Operations Manger





# Company Overview

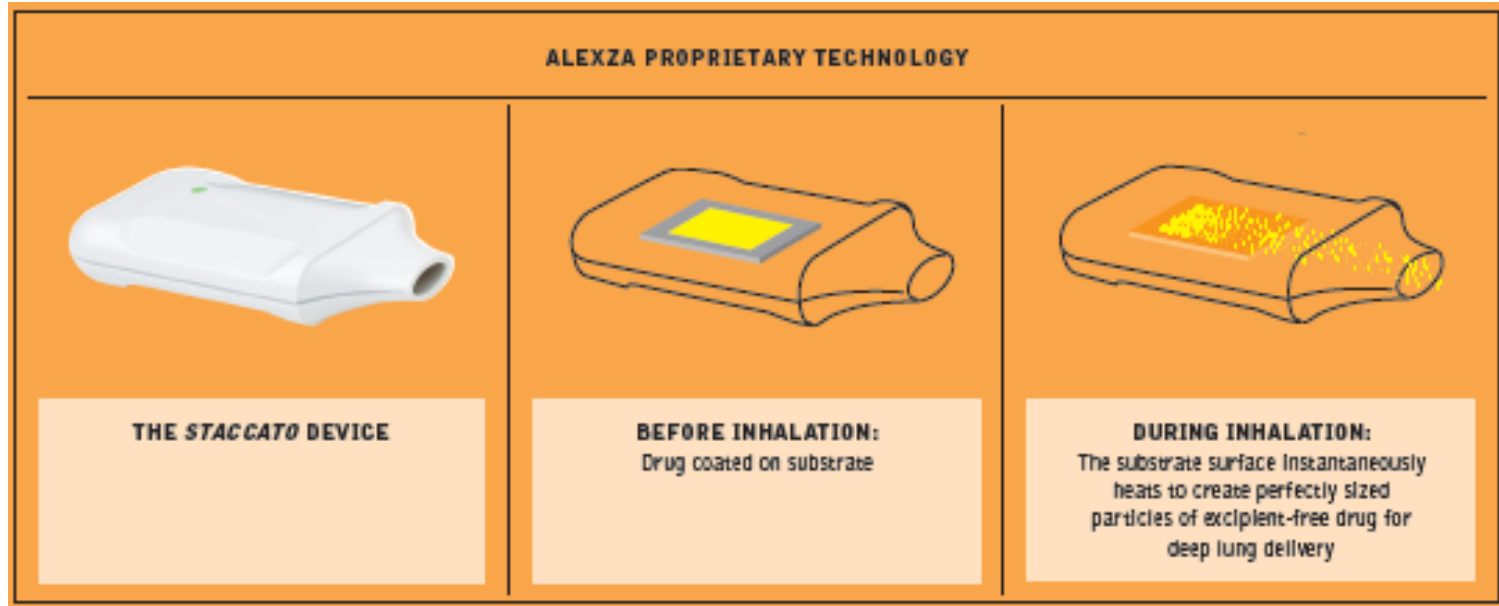
- Young, publicly traded, pre-commercial pharmaceutical company
- Focused on the treatment of Central Nervous System disorders
- *Staccato* System: vapor drug delivery
- 5 products currently in development - one has completed clinical trials

# *Staccato*<sup>®</sup> Loxapine



- Treats bipolar disorder & schizophrenia
- 8.1 million U.S. patients
- Patient self-administered – like an asthma inhaler

# Staccato System



- Alexza's Technology:
  - Single dose delivery system – heats/atomizes drug
  - Drug manufacturing/coating process
- Vapor delivers AZ-004 to the deep lung
- Acts very rapidly





# What drives our IT requirements?

- Governed by multiple regulatory agencies (e.g. FDA)
- Complex application processes
- Required to document *everything* and keep data for long time periods
- Upcoming manufacturing and QA requirements
- Like everybody else... cost!
- Extremely limited data center space

# Alexza's IT Team

- Small team of 6 supports applications, networking, storage, servers, desktops... *everything*
- Must wear many hats to cover all points
- Technology requirements must match expertise and experience
- Complex pool of applications to support: from ERP to PLM to Chromatography Data Management Systems

# Key IT Facts:


## This is our entire datacenter!





# Key IT Facts

- Just 4½ racks for everything – servers, networking, storage
- Highly virtualized environment – out of necessity
- Use well-known vendor for primary storage – great, but expensive; already at 2 racks and growing





# Challenges

- No space for growth
- Tight budgets – real cost of local storage calculated at \$8-9 per usable GB
- Completing backups on ever growing data volumes (and cost/space/IT staff time involved)
- Takes too long to provision new storage
- Don't want “off the grid” storage by rogue users



# Challenges

- Expectation of high volume data creation coming soon
- Each manufacturing line can create TBs of new data per year that we must keep per regulations




# Options and Opportunities

- Solutions considered:
  - High density disk on existing storage array
  - Secondary, lower cost storage array
  - Cloud storage
  - Deduplication and offload to another datacenter
  - Combinations of the above



# What We Chose

- Hybrid Cloud Storage Solution
- Amazon S3 cloud accessed via on-premise appliance: Cirtas Bluejet Cloud Storage Controller
- Looks like an iSCSI SAN – zero learning curve
- 2U – unlimited capacity



# Concept to Reality

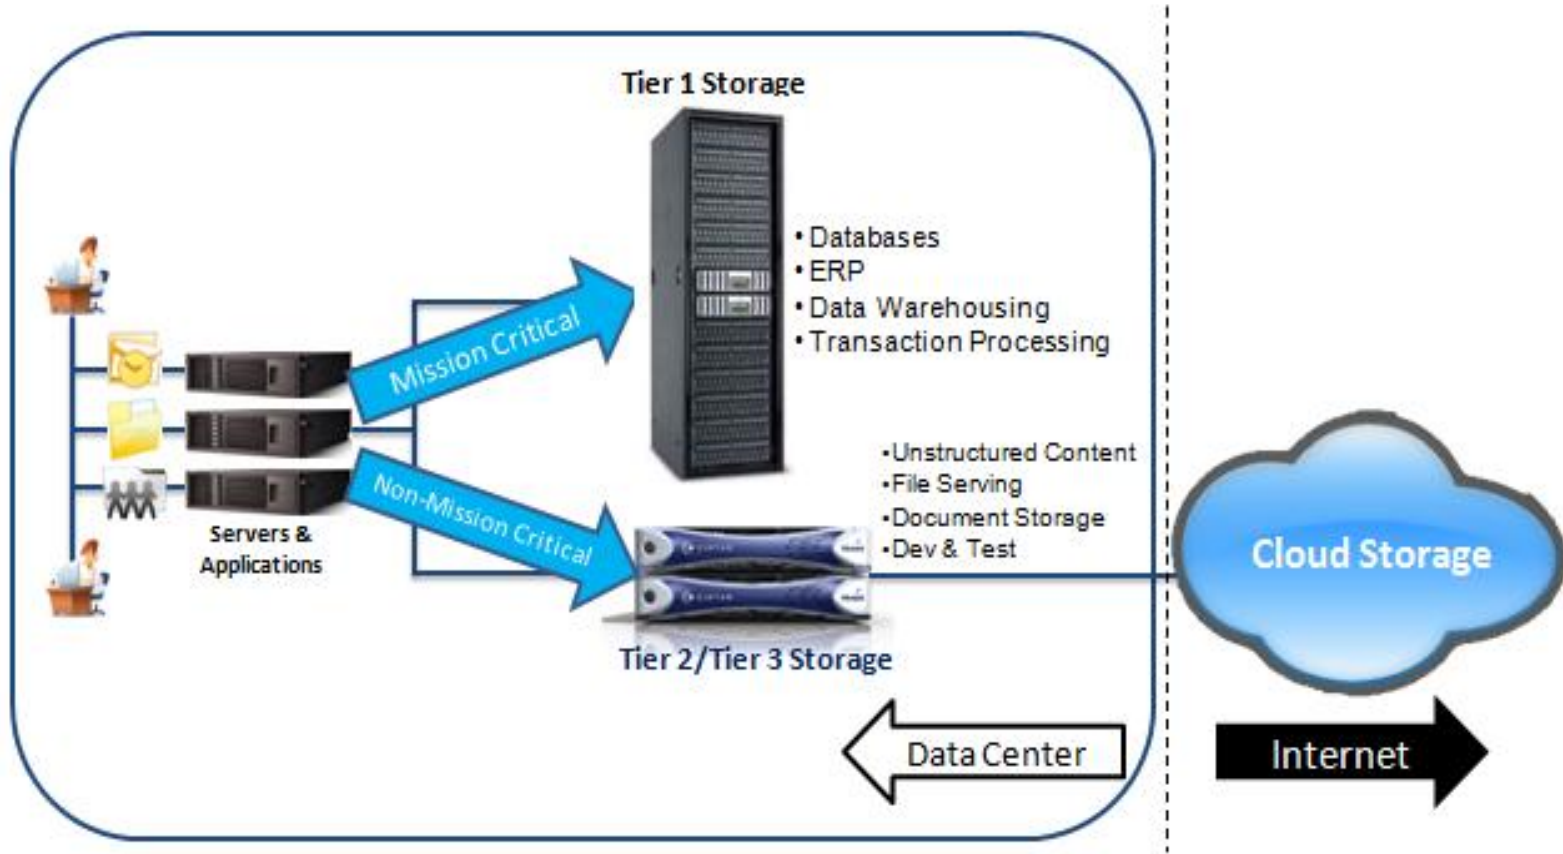
## ■ Pros

- Up and running in less than an hour
- No application changes – instantly compatible
- Immediate access to storage as needs arrive
- Local caching improves performance compared to pure cloud approach while lowering data transfer fees
- Deduplication lowers static costs of cloud storage and improves bandwidth utilization
- We can dedicate our existing storage to Tier1 workloads – everything else goes to the cloud

## ■ Cons

- WAN dependent – currently uses our 10Mbps link
- Does not offset IOPS intensive needs best suited for primary storage

# Our Environment Today



- Tier 1 stays in on-premise storage array
- Tiers 2 and 3 offload to the cloud

# It works!

- ✓ Growth pressure on our existing NAS has slowed – Tier 1 disk purchases averted
- ✓ Flexibility of storage has saved recovery time (moving unused VMs out to the cloud before final deletion)
- ✓ We can respond quickly to end user needs (“just in time” storage)
- ✓ No need to backup data with our Bluejet Cloud Storage Controller – it takes care of data protection and disaster recovery on its own





# Why Transition to Hybrid Cloud Approach?

- Data Discrimination: Not everything deserves the highest speed spindles
- Many things just have to be accessible when needed – a few seconds is OK
- Newton's First Law of Motion tends to apply to data
  - Data at rest tends to stay at rest
  - Hybrid cloud is ideal for dormant and “semi-active” data
- Vastly lowered costs and management requirements

# Results

- ✓ Achieved deduplication ratios of 10:1 – our cost to store in the cloud is thus 1/10<sup>th</sup> what you would expect!
- ✓ Thin provisioned 30TB, currently using 3TB
- ✓ Have offloaded 387GB of data to the cloud (3TB of data with 10:1 deduplication)
- ✓ Amazingly low Amazon S3 bill: was \$35.27 last month - equates to \$0.42/GB over 3 years (compared to \$8-9/GB onsite). Cloud Storage Controller pays for itself quickly.
- ✓ Completely transparent to end users when using the appropriate data sets

# Next steps

- **Expand usage**
  - Move more non-business critical data
- **Expand WAN size efficiently**
  - Monitor and move only when needed
- **Explore NDMP receiver VM** for disk to disk backup of primary storage



# Cheating Death in the Datacenter...

## Questions?

David A Jones  
IT Operations Manger

**ALEXZA**  
PHARMACEUTICALS