



**Driving Innovation
Through the Information
Infrastructure**

SPRING 2011



Make Your Existing File Storage Infrastructure “Cloud-Ready”

Richard Gillett

VP, Data Systems Architecture

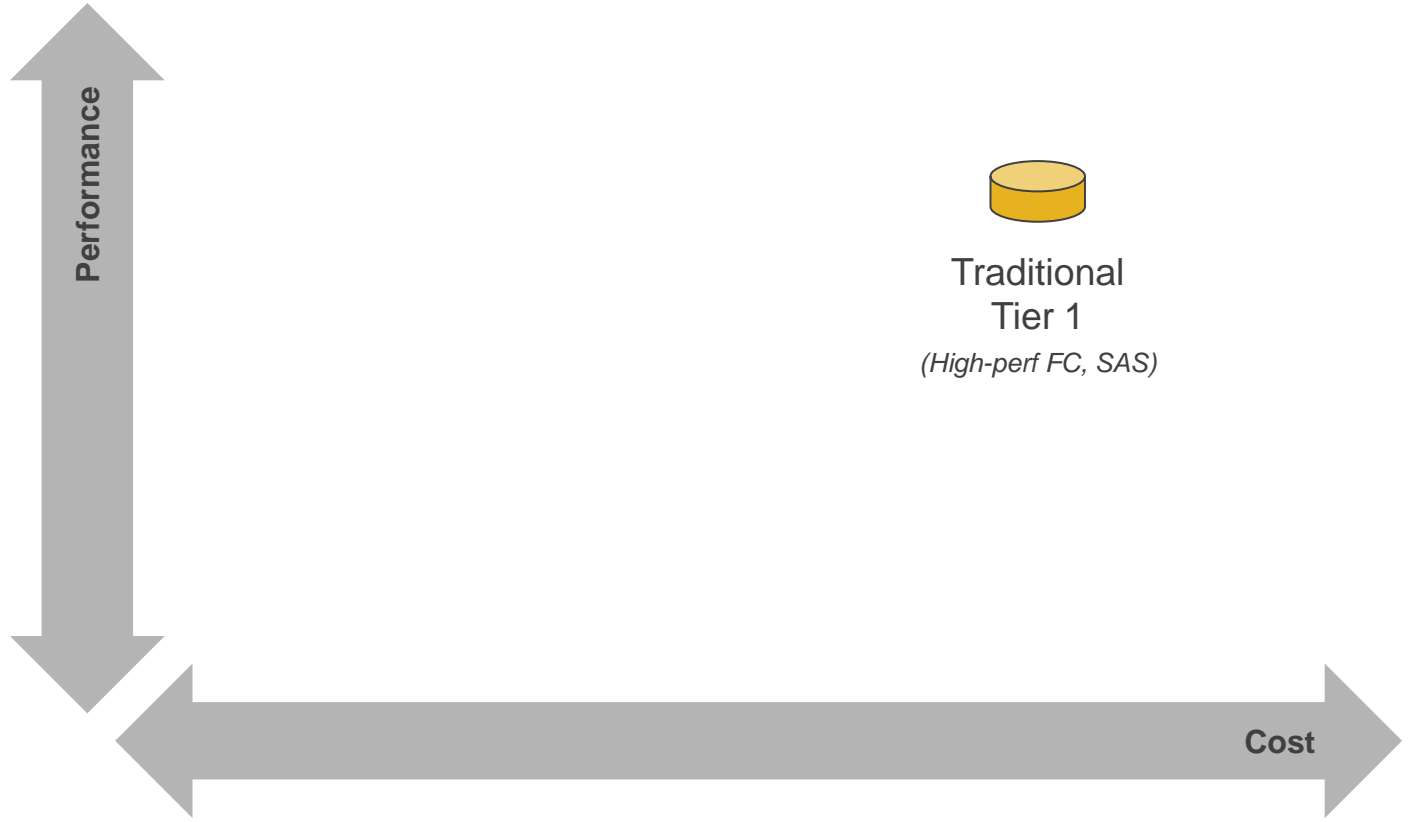
F5 Networks



Agenda

- What is Cloud Storage?
- Technical Barriers to Using Cloud Storage
- Getting Ready for Cloud
 - Options for integration
- Advanced Applications
 - Integrated archive example
- Summary

Storage Continuum

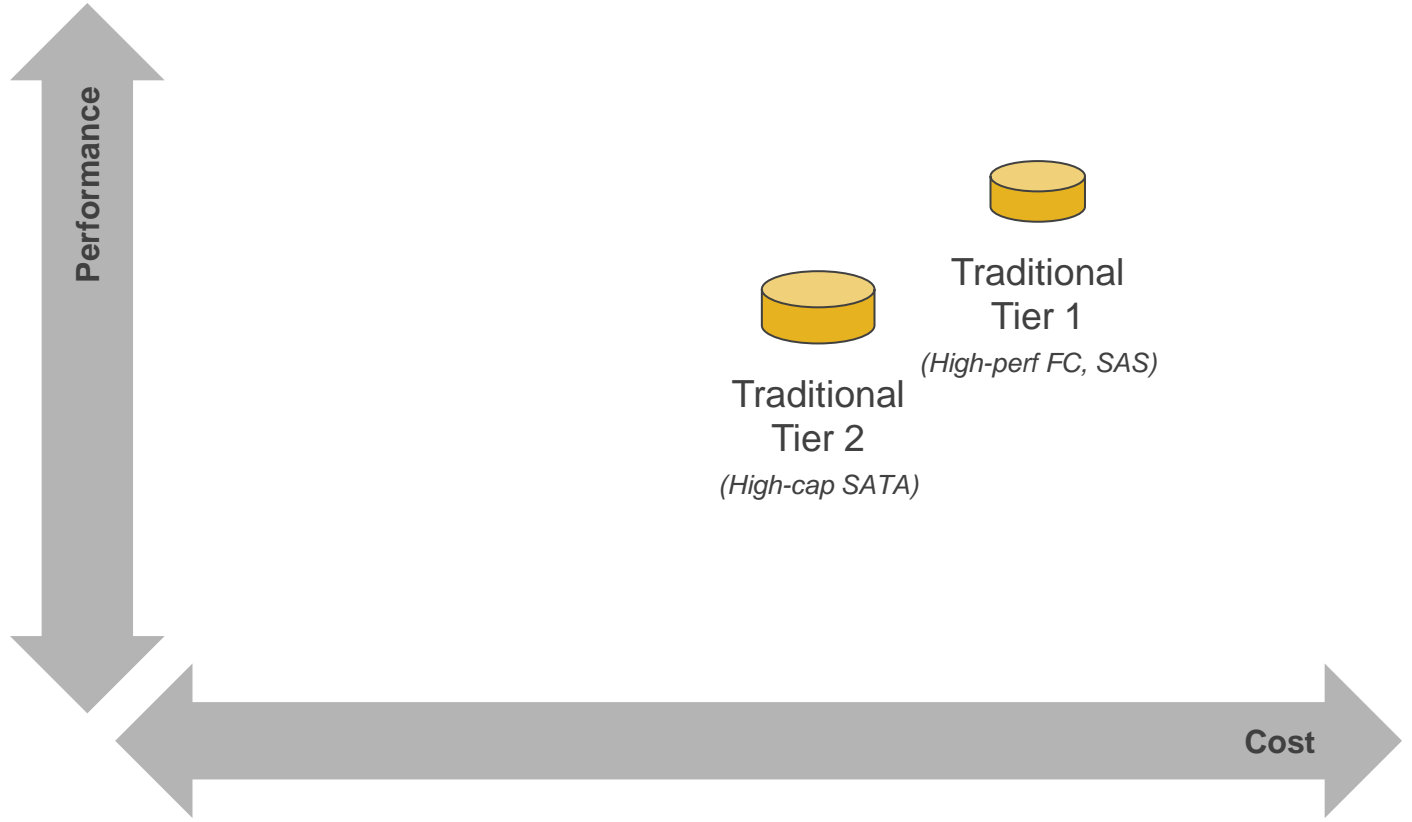


Traditional
Tier 1
(High-perf FC, SAS)

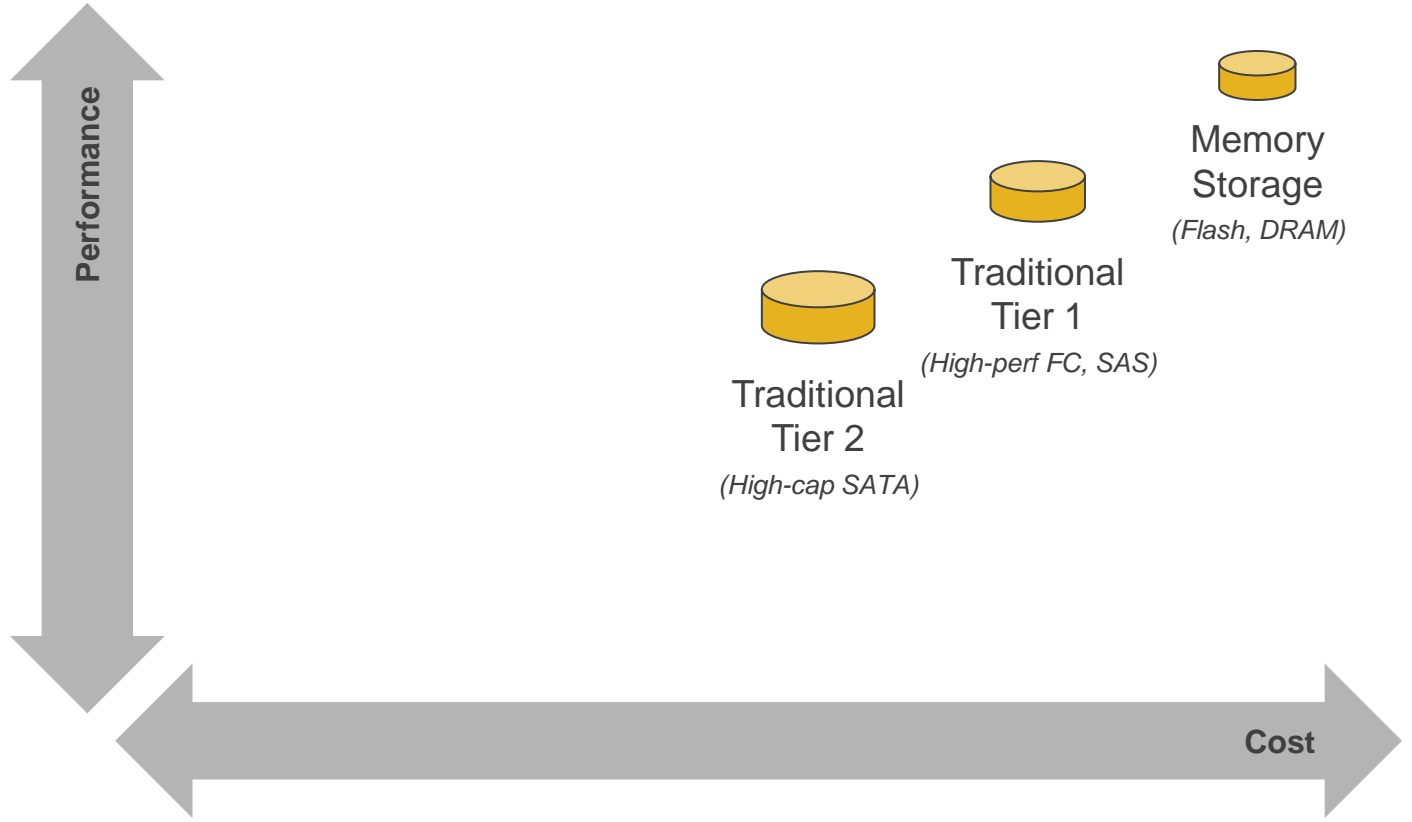
Performance

Cost

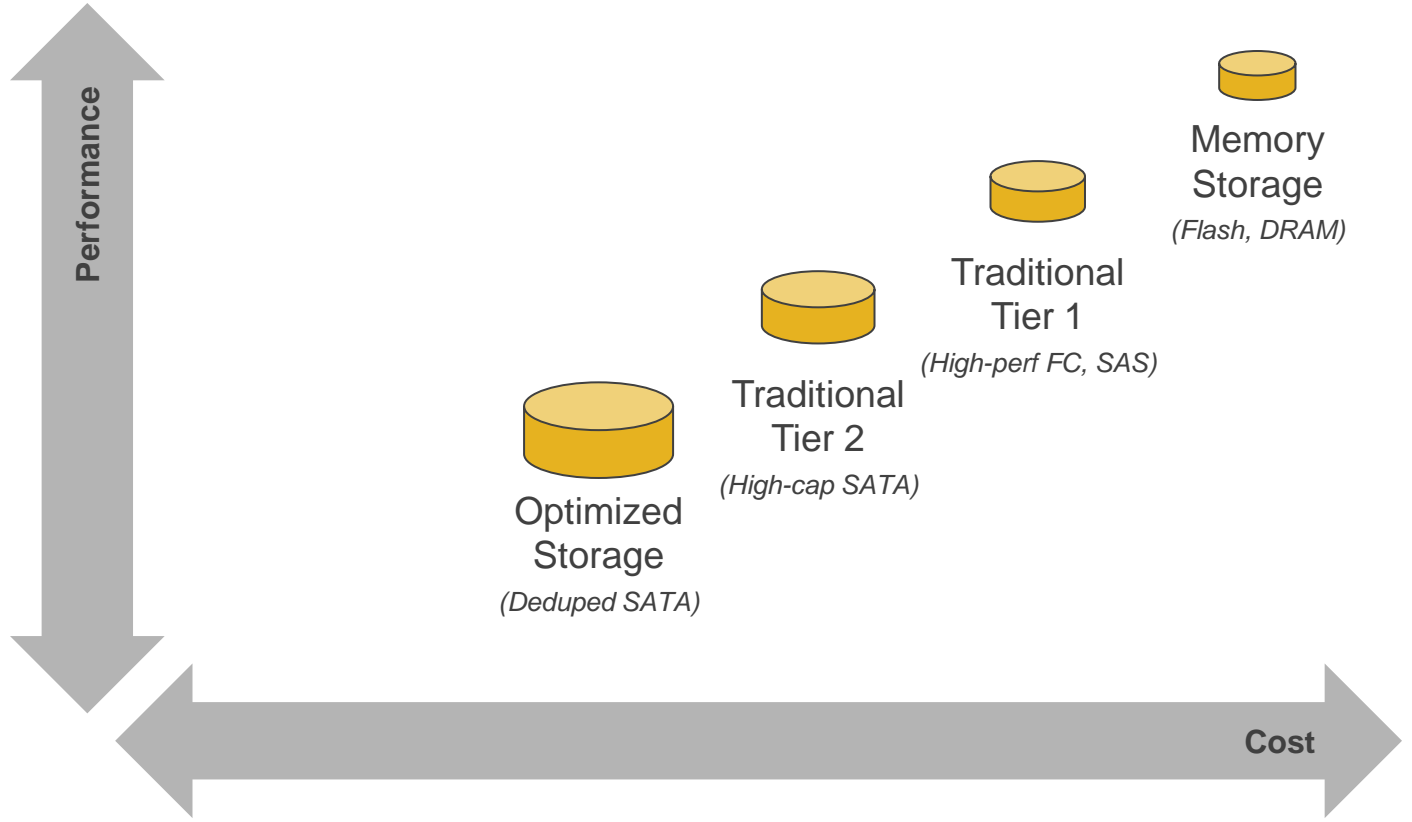
Storage Continuum




Storage Continuum



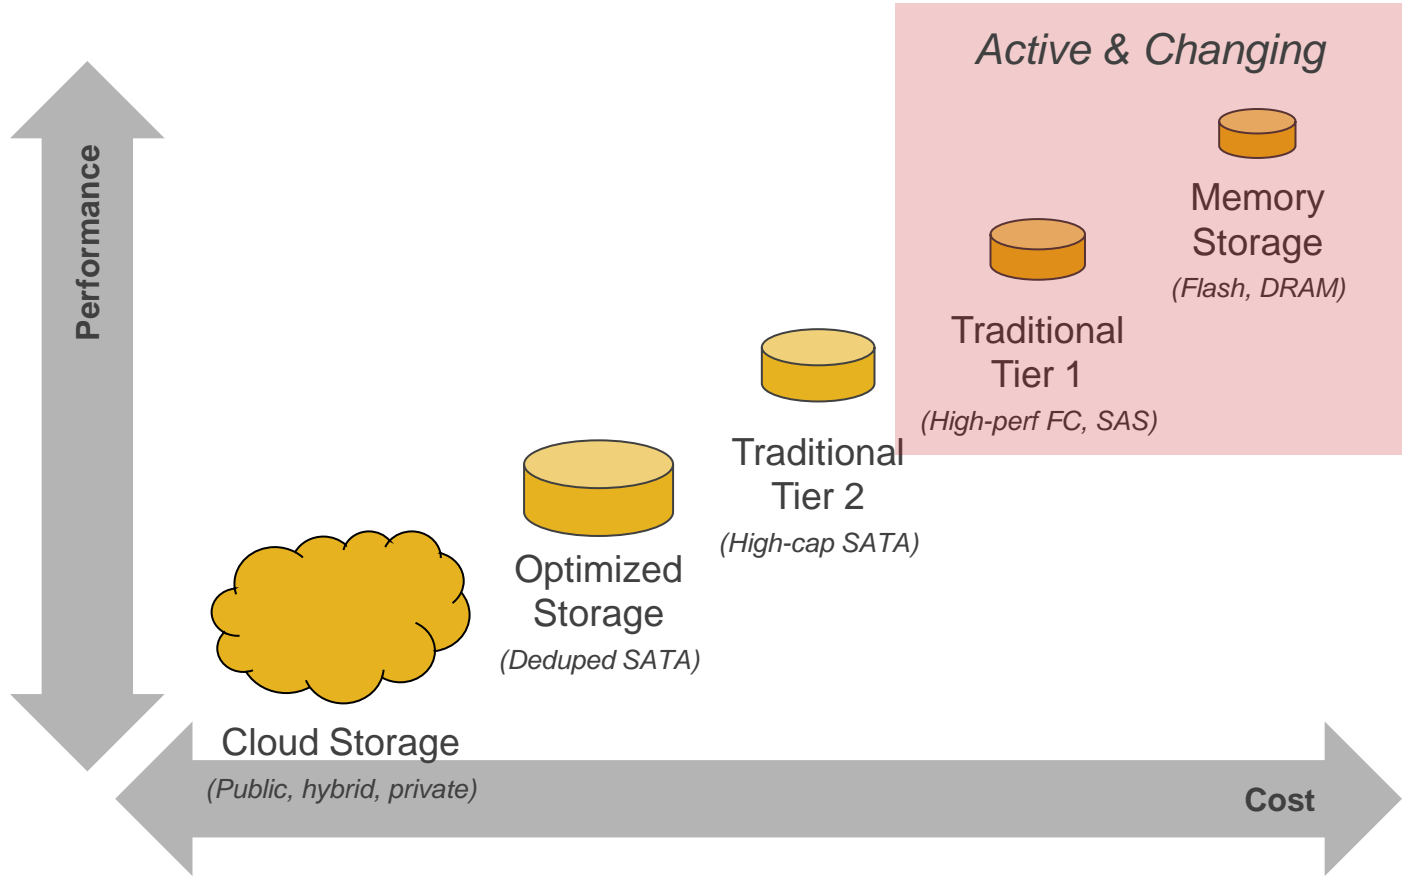
Storage Continuum



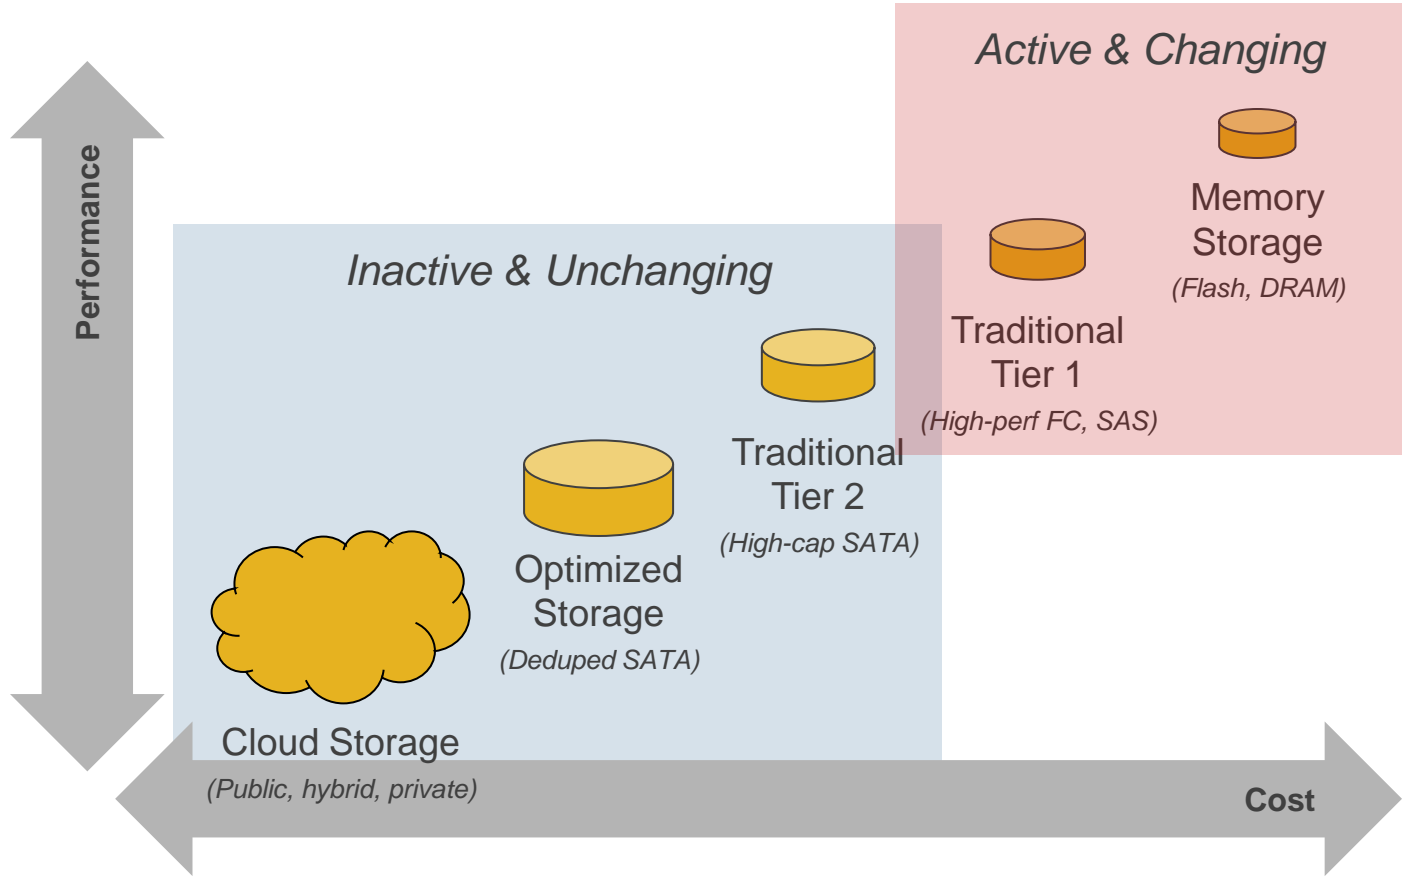
Storage Continuum



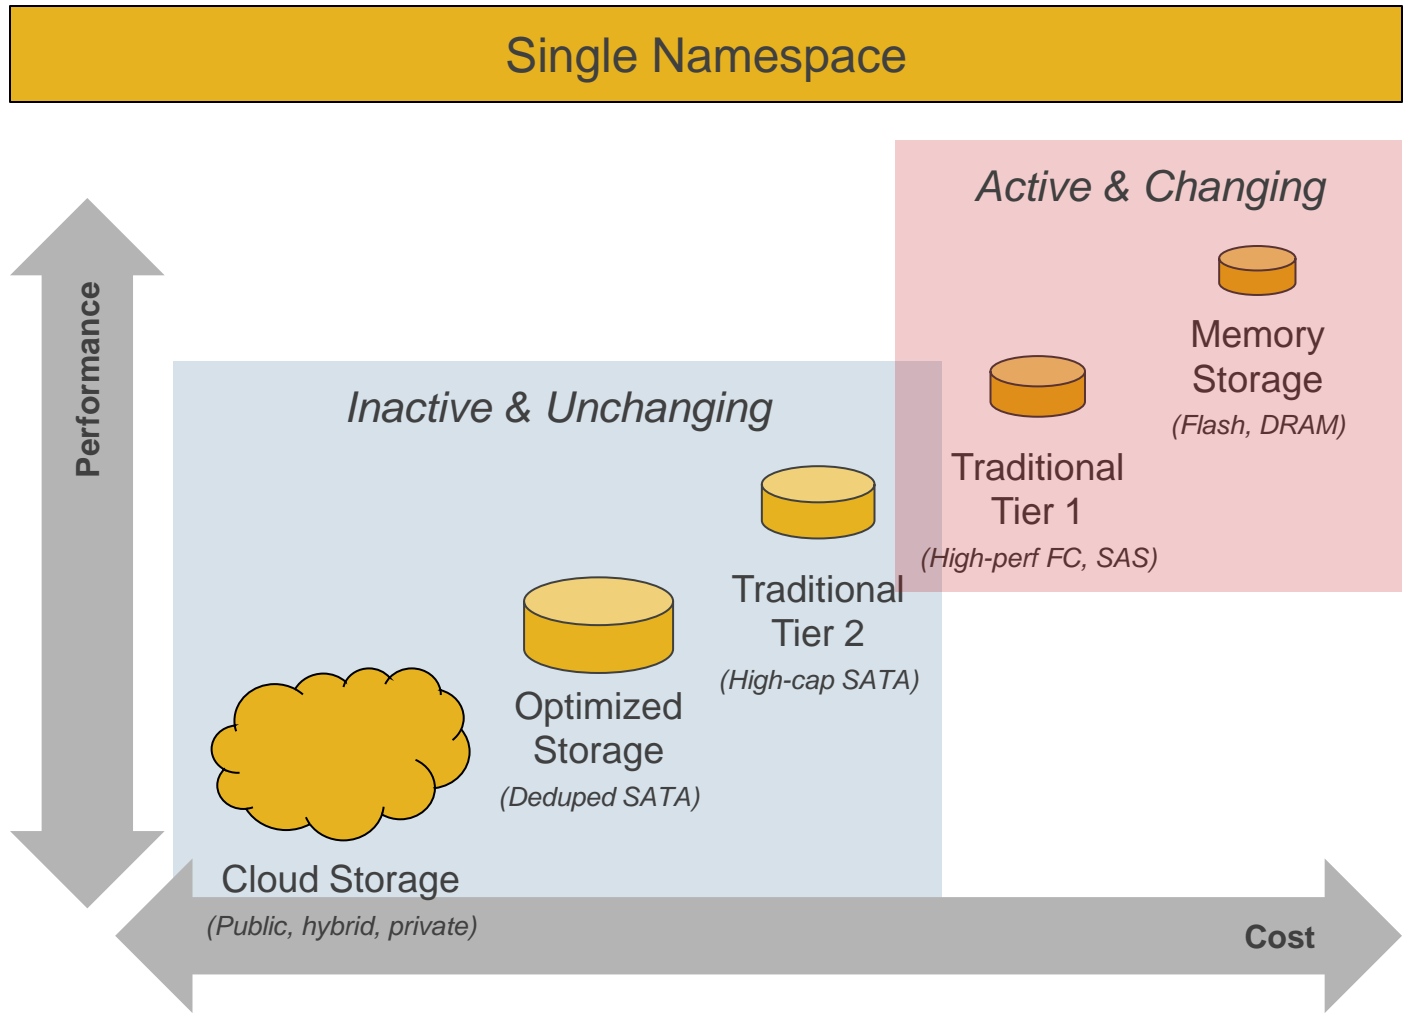
Big Picture Opportunity



Big Picture Opportunity



Big Picture Opportunity




Comparing Performance

- General guidelines:

Type of Storage	Aggregate Bandwidth	Single Operation Latency
Local Tier 0 (SSD)	> 1000 MB/s	< 500 μ s
Local Tier 1 (SAS)	500 MB/s	0.5 to 5 ms
Local Tier 2 (SATA)	100 MB/s	.5 to 15 ms
Cloud	.1-10 MB/s	10 ms to >10s

Strong Interest In Cloud

Would you consider private cloud as an alternative to traditional NAS Storage?



Source: F5 Networks, July 2010

What Is Cloud Storage ?

- No consistent definition
- For our purposes today, cloud storage is:
 - Remotely accessed
 - Accessed via a WAN
 - Optimized for long-term retention
 - Typically implemented as some form of object store
 - Not “natively” presented as a file system
 - Typically a web-based API
 - Might be “rented” (bought as a service) or “owned” (deployed conventionally)
 - More on the next slide

Simple Cost Comparison

Characteristic	Own Tier-1	Own Tier-2	Rent Cloud
Purchase price (\$/GB)	\$10	\$2	-
Rental cost (\$/GB-Month)	-	-	\$0.15
HW/SW Maintenance (% of price over 36 months)	15% per year	15% per year	-
Power and cooling costs (\$/GB over 36 months)	\$1.8	\$0.2	-
Effective Cost Per GB Over 36 Months 100% utilization; 0% access	\$16.30	\$3.10	\$5.40
Effective Cost Per GB Over 36 Months 50% utilization; 0% access	\$32.60	\$6.20	\$5.40

Own or Rent – A Simple Model


- Profile for renting storage
 - Low storage utilization today
 - Modest total capacity
 - Temporary peaks in storage demand
- Profile for owning storage
 - High storage utilization today
 - High total capacity
 - Willing to deploy cloud-like technology in place of traditional NAS
 - Storage capacity demand continually grows



CHALLENGES

Goals and Constraints

Users



What users care about:

- No change to how they access files
- Having a single location for accessing files
- Maintaining a reasonable level of performance

IT



What IT cares about:

- Managing growth
- Improving utilization
- Reducing costs
- Reducing overall complexity
- Ensuring data security
- Maintaining data availability
- Meeting regulatory compliance requirements



Challenges and Solutions

Challenges

Maintain transparent access to files

Move appropriate files to and from cloud

Remote storage latency

Data security

Different “storage model”


Challenges and Solutions

Challenges	Solutions
Maintain transparent access to files	<ul style="list-style-type: none"> • Data redirection <u>within</u> the file server • File virtualization <u>above</u> the file server
Move appropriate files to and from cloud	<ul style="list-style-type: none"> • Automatic classification + policy + migration • Manual
Remote storage latency	<ul style="list-style-type: none"> • Only store “inactive” data in cloud • Locally cache all file metadata • Locally cache all file data
Data security	<ul style="list-style-type: none"> • 100% data and metadata encryption before leaving customer site
Different “storage model”	<ul style="list-style-type: none"> • File system gateway to cloud API




INTEGRATION

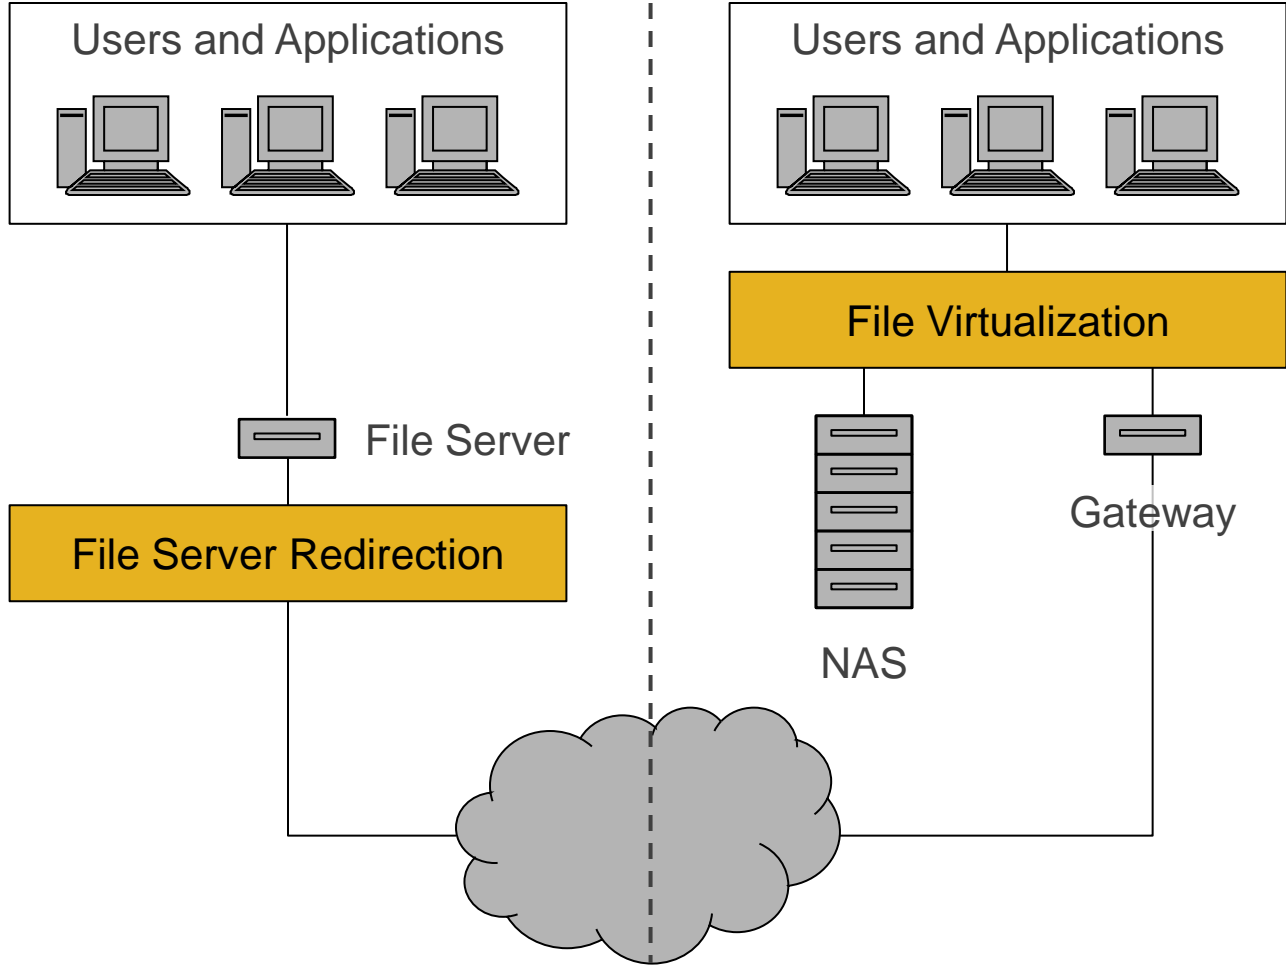
File Server Redirection



File Virtualization




Two Options









File Server Redirection




File Server Redirection







File Server Redirection

-  Foo.txt file (metadata + data)
-  Foo.txt (metadata & pointer to data)
-  Foo.txt data







File Server Redirection

-  Foo.txt file (metadata + data)
-  Foo.txt (metadata & pointer to data)
-  Foo.txt data




File Server Redirection


-  Foo.txt file (metadata + data)
-  Foo.txt (metadata & pointer to data)
-  Foo.txt data



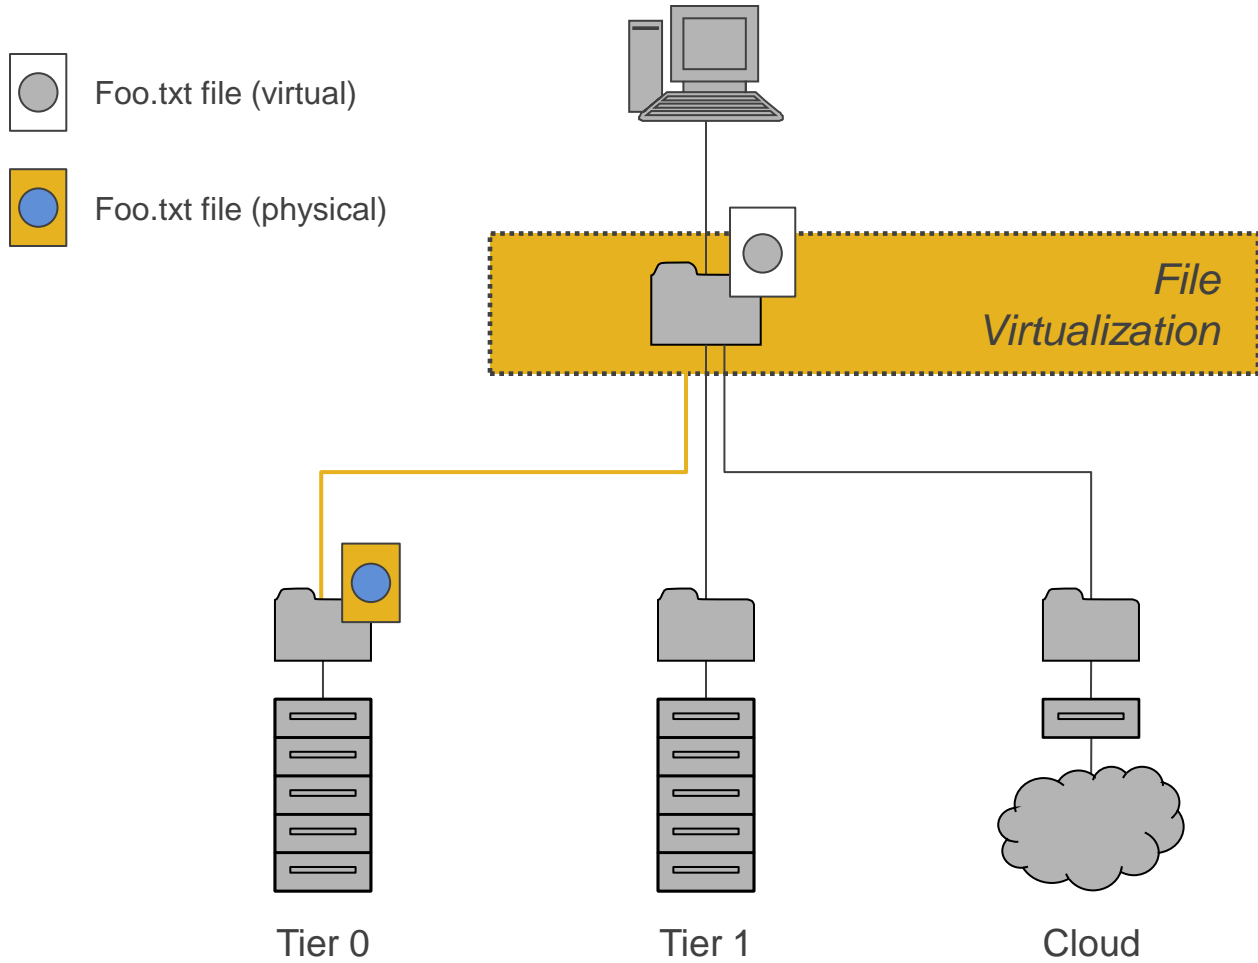
File Virtualization



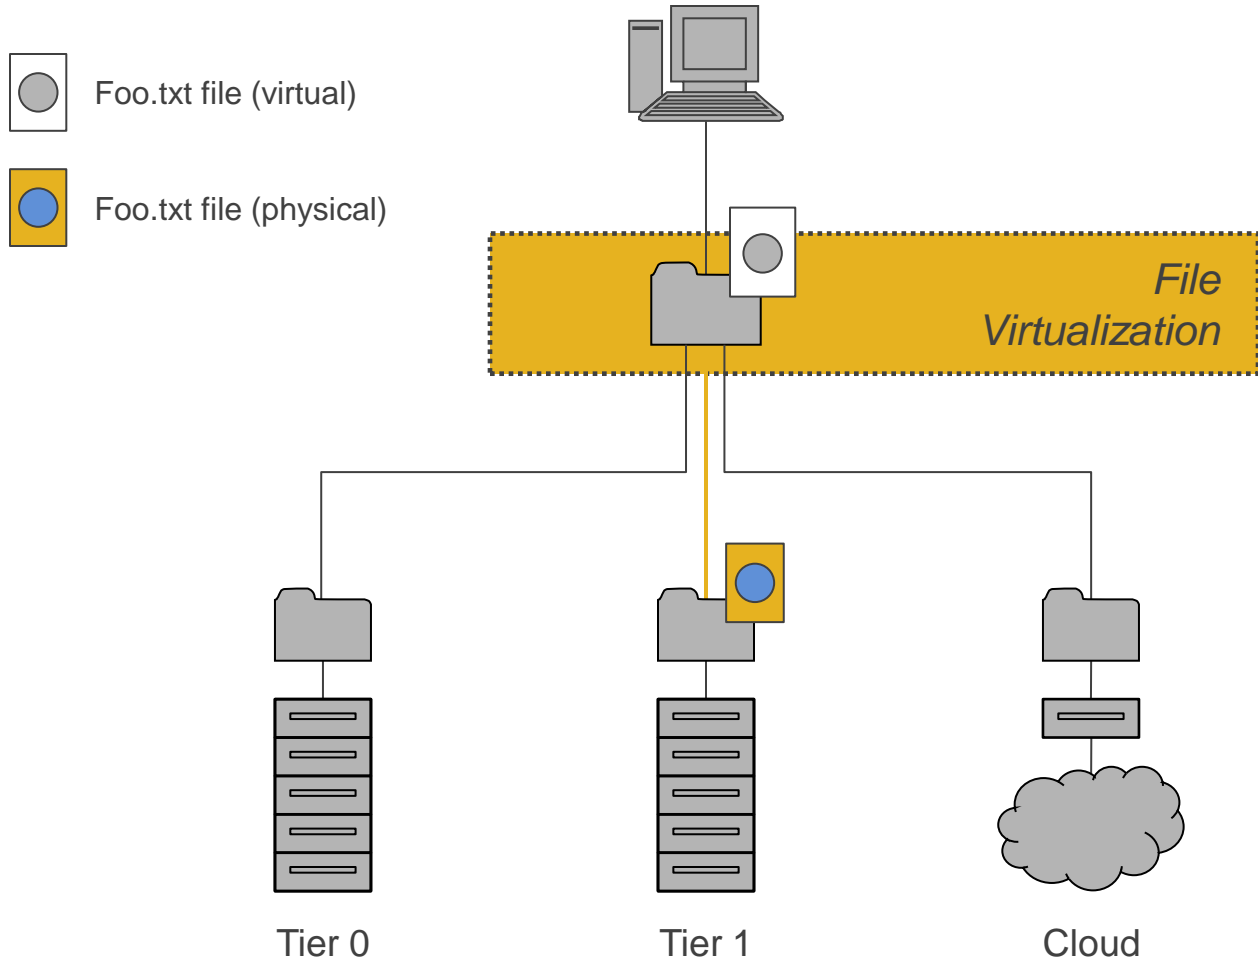
File Virtualization




File Virtualization



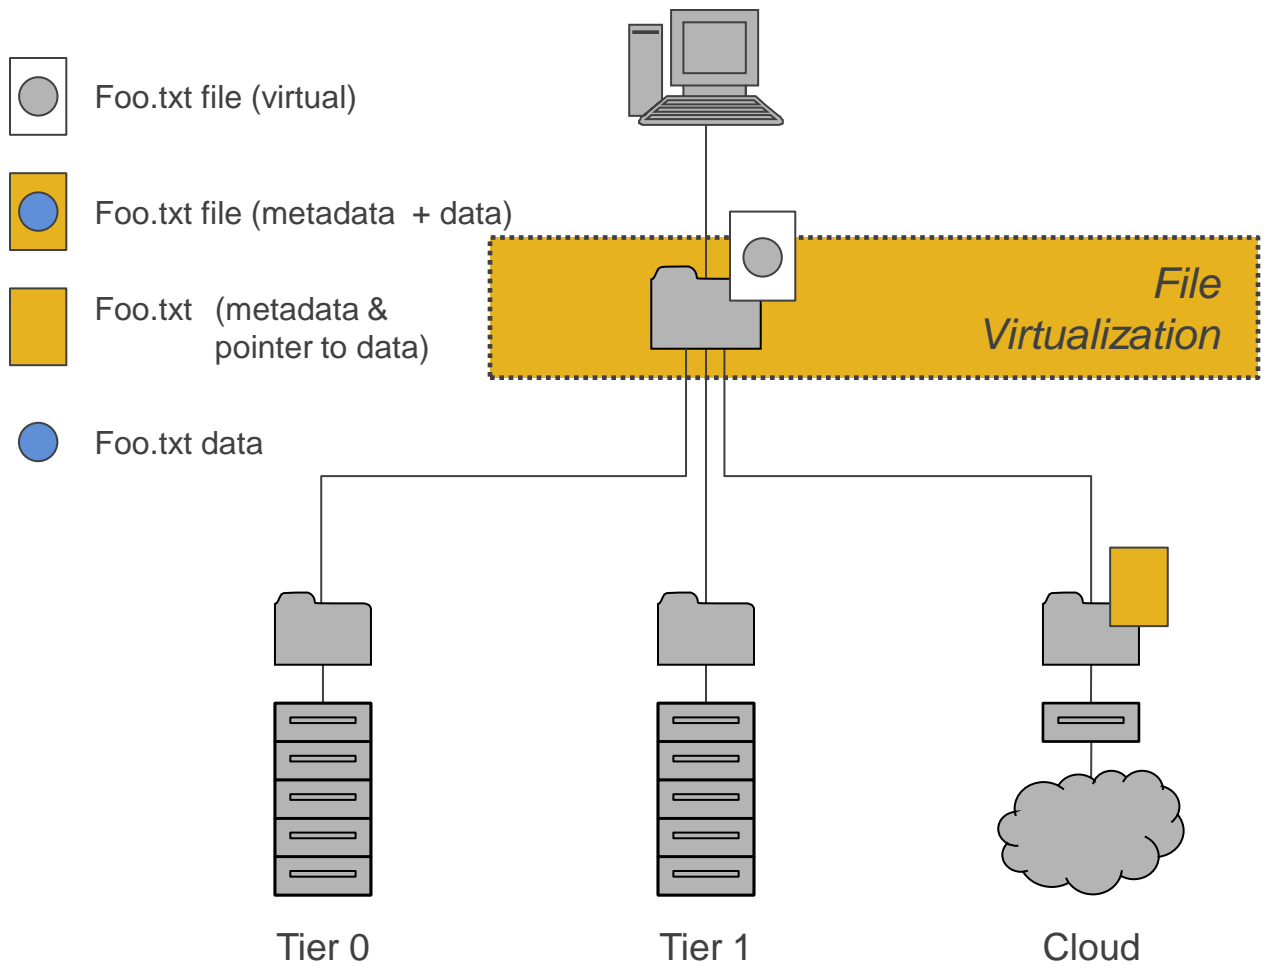
File Virtualization



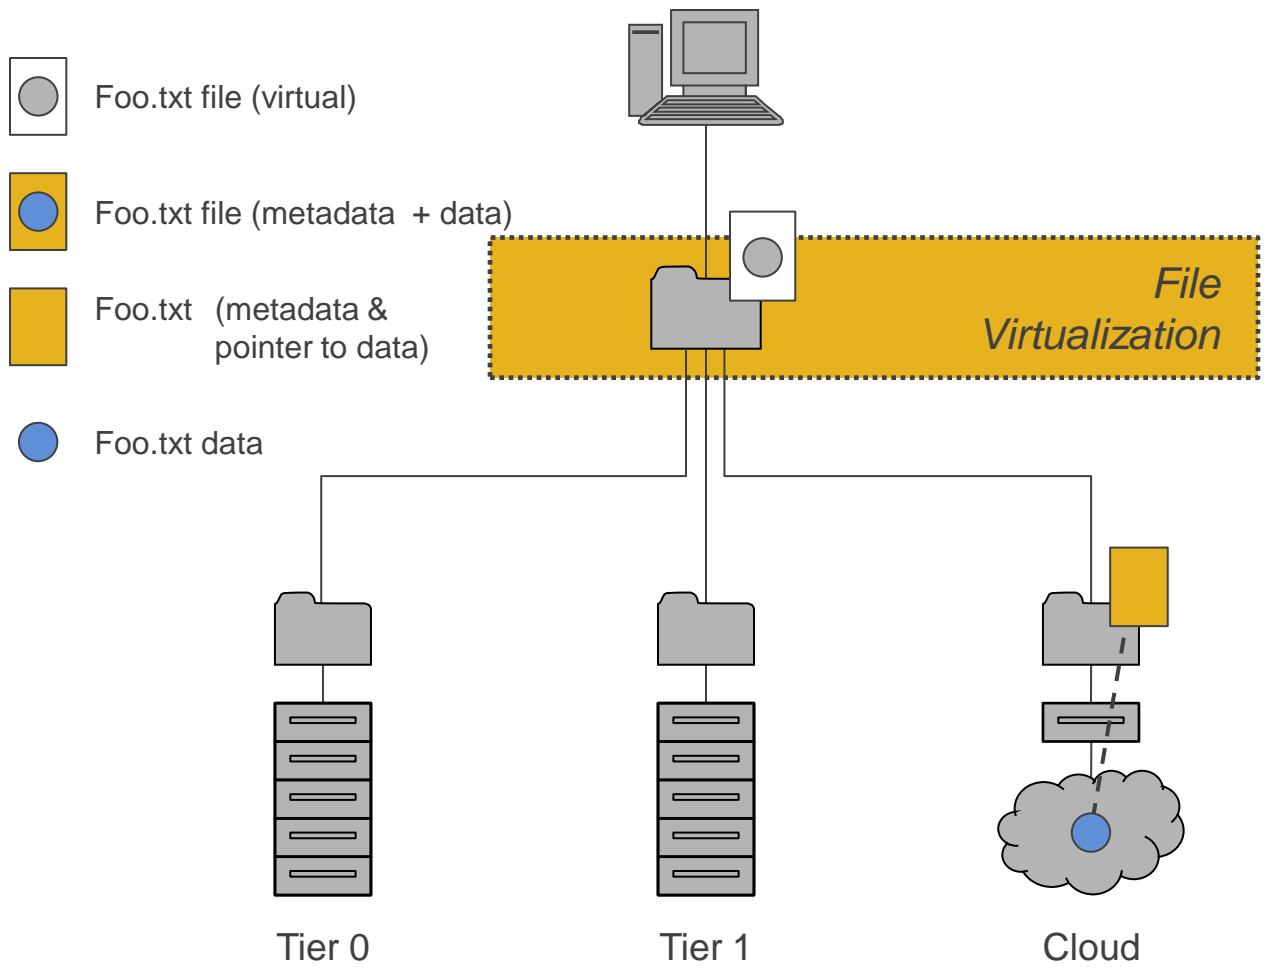
File Virtualization



File Virtualization



File Virtualization



Comparison of Approaches

Aspect	File Server Redirection	File Virtualization
Availability of software-only options (VM)	Yes	Yes
Scale	Limited by single file server	Limited by file virtualization device
Choice of cloud technology	Limited by file server support	Any
Advanced cloud services	Limited by file server support	Any



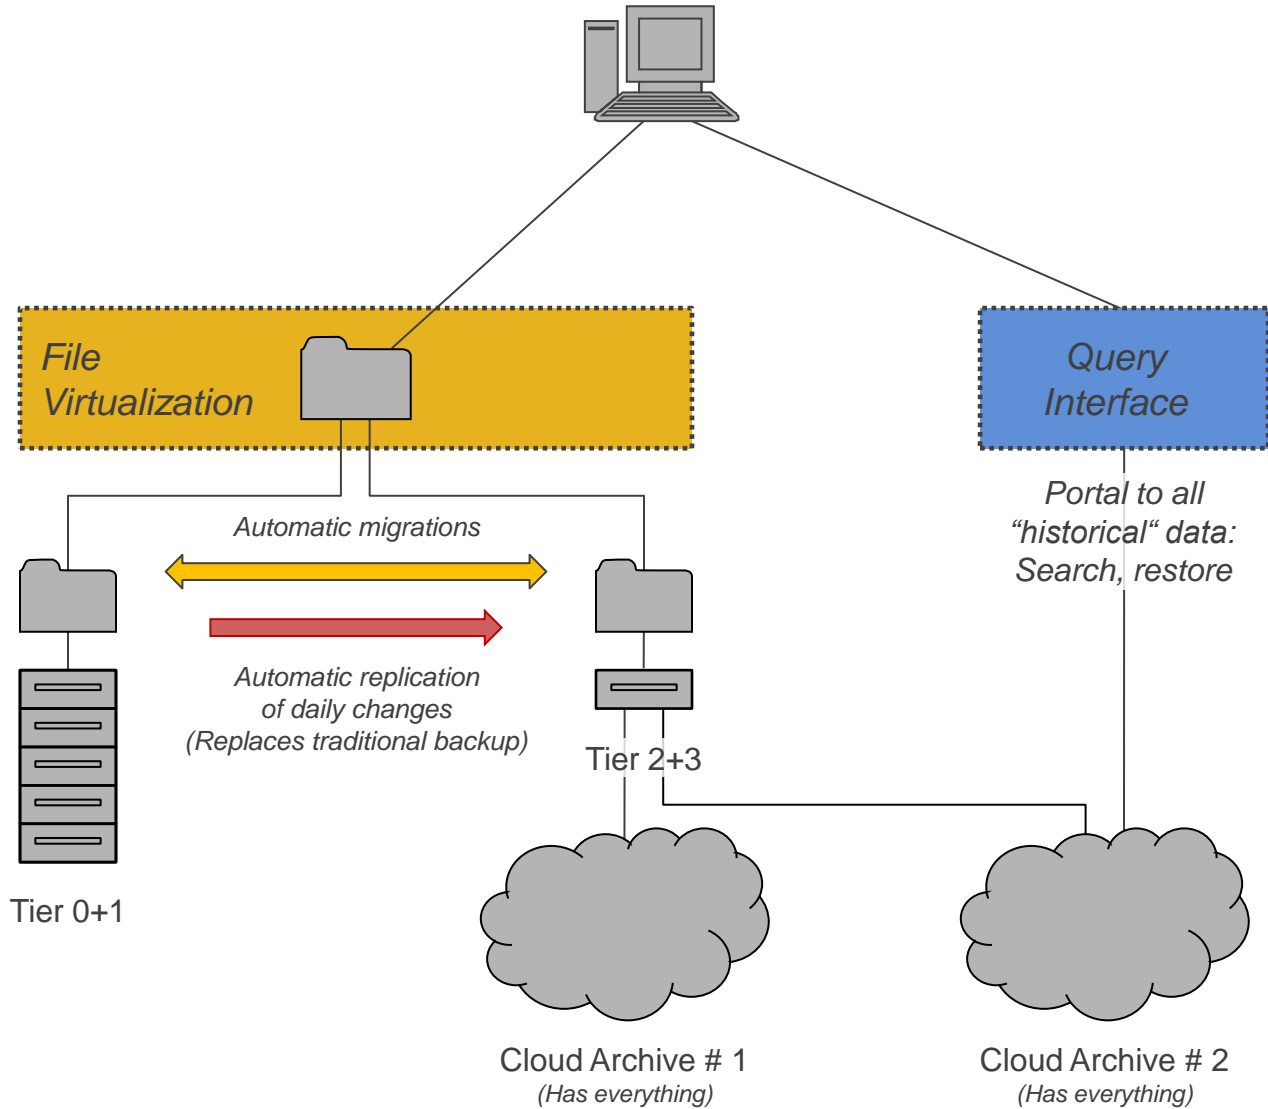
ADVANCED APPLICATIONS

Integrated (Cloud-Based) Archive


- What is it?
 - A unified approach to data management that integrates and optimizes tiering, backup and recovery, discovery, and compliance
- Simple example:
 - In an integrated archive, a single copy of data* in the cloud archive can be used simultaneously as the data storage for multiple applications:
 - Tier 2
 - Backup and recovery
 - Compliance

**Assume this “copy” is redundantly stored by the Integrated Archive*


Integrated Archive

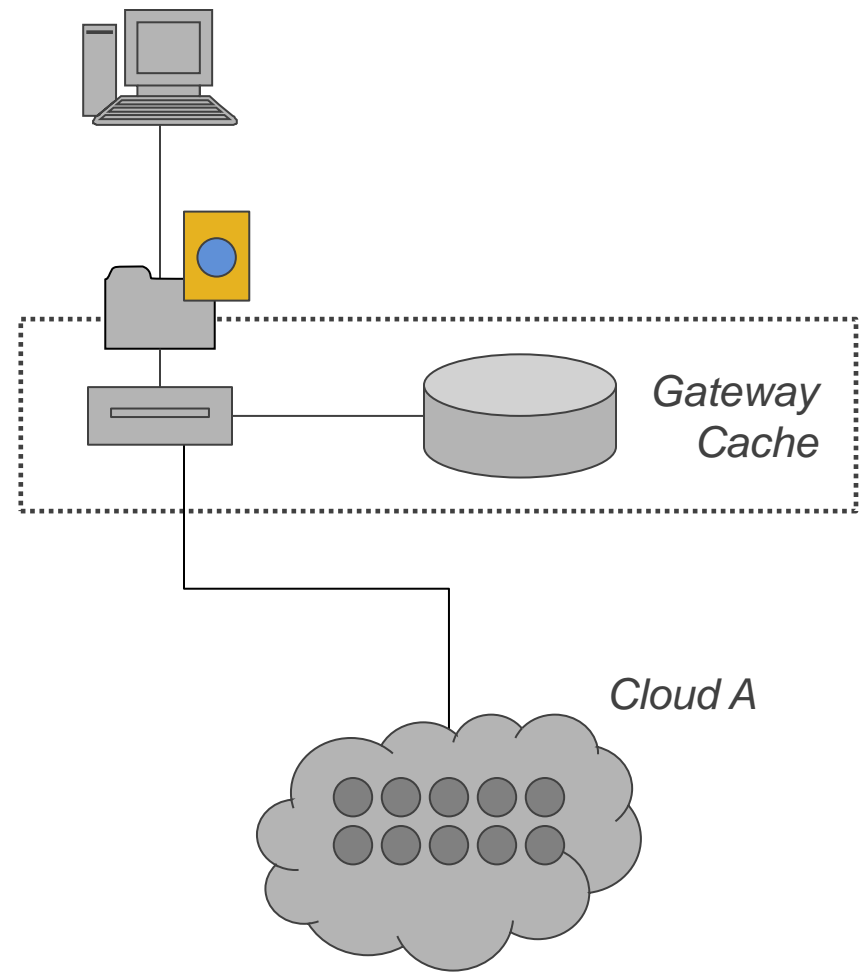


Integrated Archive In Action




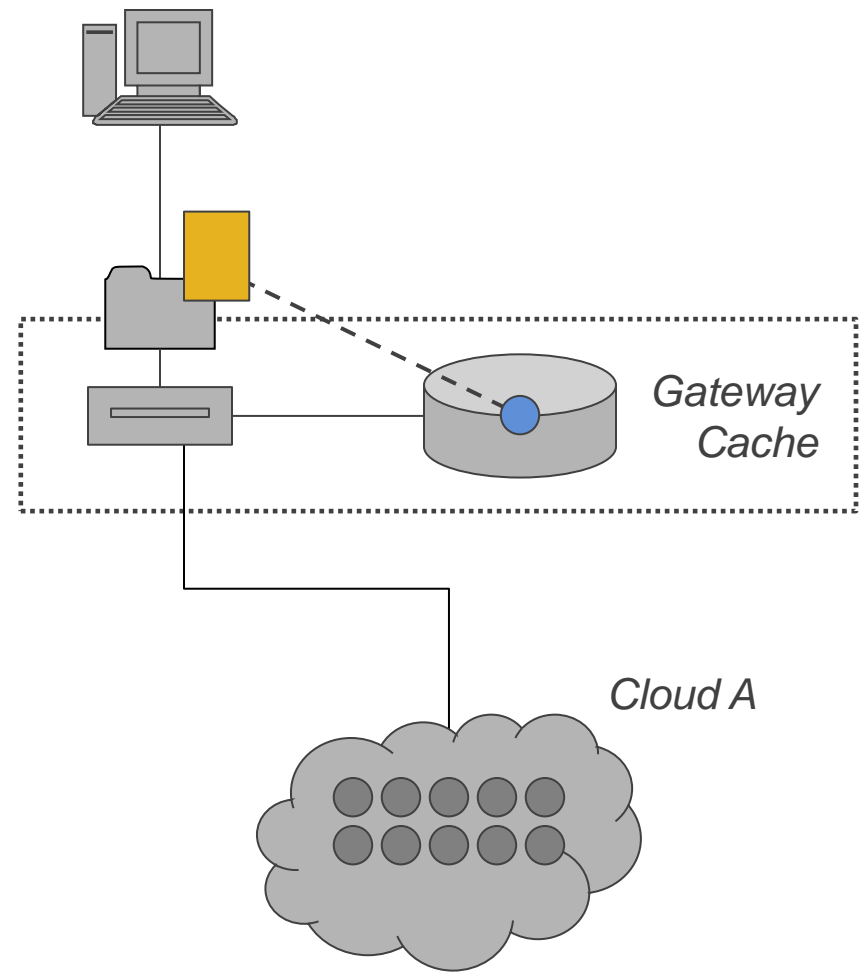
Integrated Archive In Action

 Foo.txt version 1





Integrated Archive In Action

 Foo.txt version 1






Integrated Archive In Action

 Foo.txt version 1




Integrated Archive In Action

-  Foo.txt version 1
-  Foo.txt version 2




Integrated Archive In Action

-  Foo.txt version 1
-  Foo.txt version 2

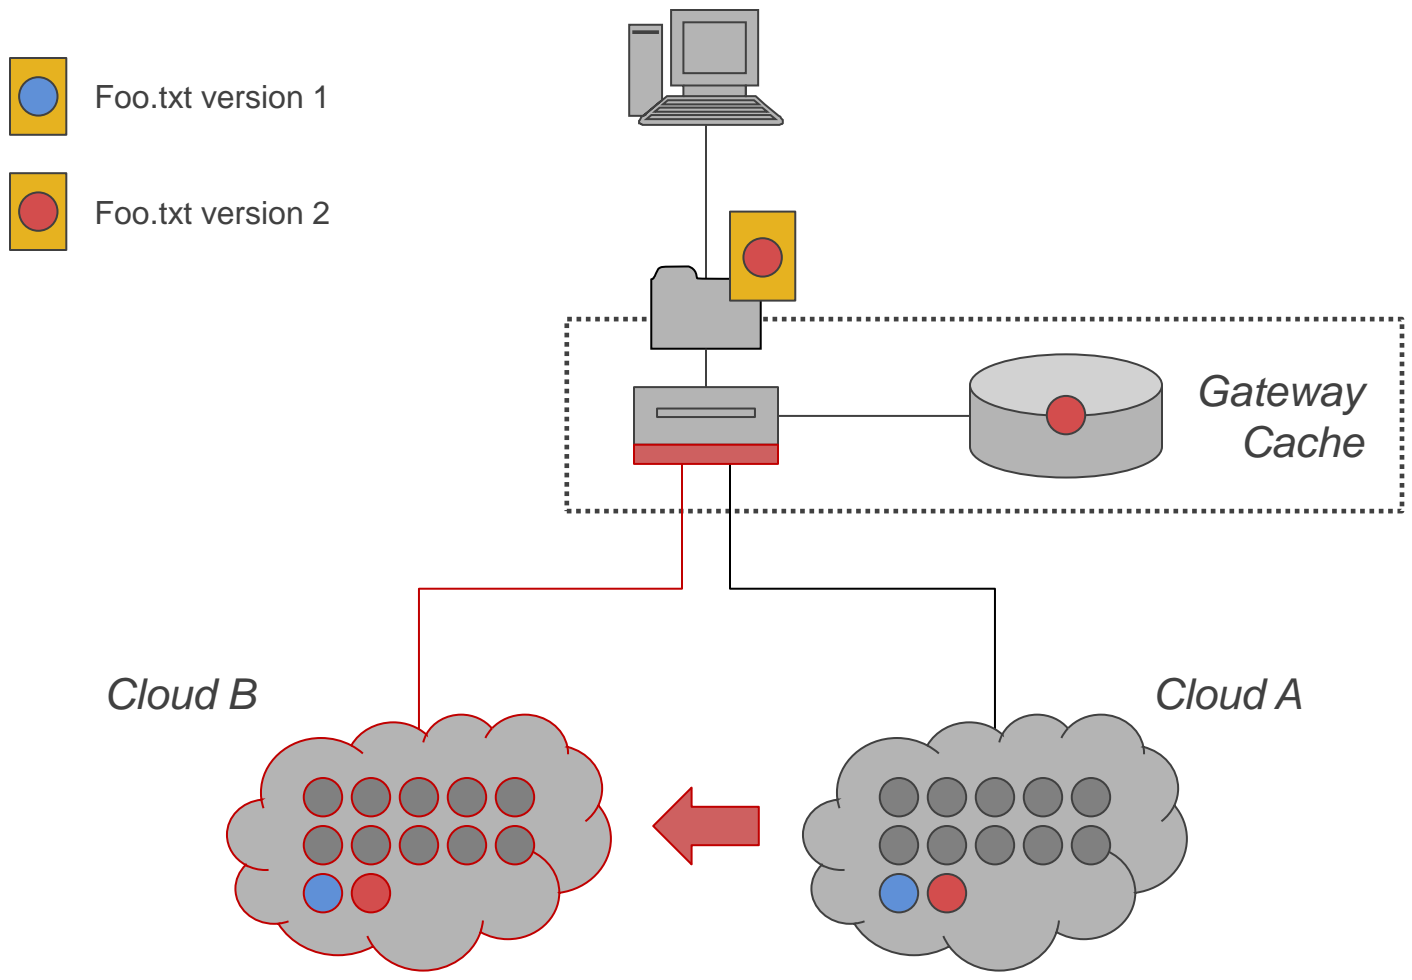


Integrated Archive In Action

-  Foo.txt version 1
-  Foo.txt version 2



Integrated Archive In Action





Summary

- Cloud storage technology is an attractive choice for long-term storage of files
 - Important to understand cloud technology
 - Independent of the options to deploy (rent/buy)
- There are two options to integrate cloud storage with existing file storage environments
 - Best option will depend on your environment and data storage requirements
- Cloud storage technology enables advanced data management
 - Example: Integrated Archive