Log Based Storage

Mark O’Connell
mark@mkoconnell.com

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Abstract – Log Based Storage

- While the concept of a log is certainly not a new concept in computer science, until recently logs have been used as a part of an application, typically for crash recovery purposes or sometimes for auditing/debugging purposes. More recently, logs have been emerging as a first class storage concept in and of themselves, being used in distributed environments as a mechanism for communication, as a mechanism of persistence and recovery for services, and as an enabler for query optimized data structures in complex systems. This tutorial will cover the history of log storage, starting with its use in databases and transaction logs, will contrast log storage vs block, file, and object storage, and then will examine the role of log storage in distributed systems and microservice environments via event logs and CQRS patterns (Command Query Responsibility Separation).

What is a log?

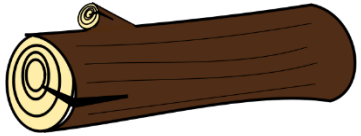

- Many real world examples
  - Visitor log, access log, arrival log, system log
- Characteristics
  - Ordered set of events
  - Readable by many
  - Append only
  - Reads from anywhere, though characteristics differ

Traditional uses

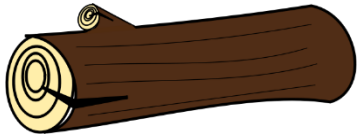

- Audit log – what happened?
- Debugging logs – what went wrong?
- Recovery log
  - DB
  - FS
- Source code control change log – who broke it?

What's changing?

- Analytics
- Real time streams
- Replayability
- Distributed / Decoupled Systems
- Microservices

Consequences

- Logs become a communication mechanism

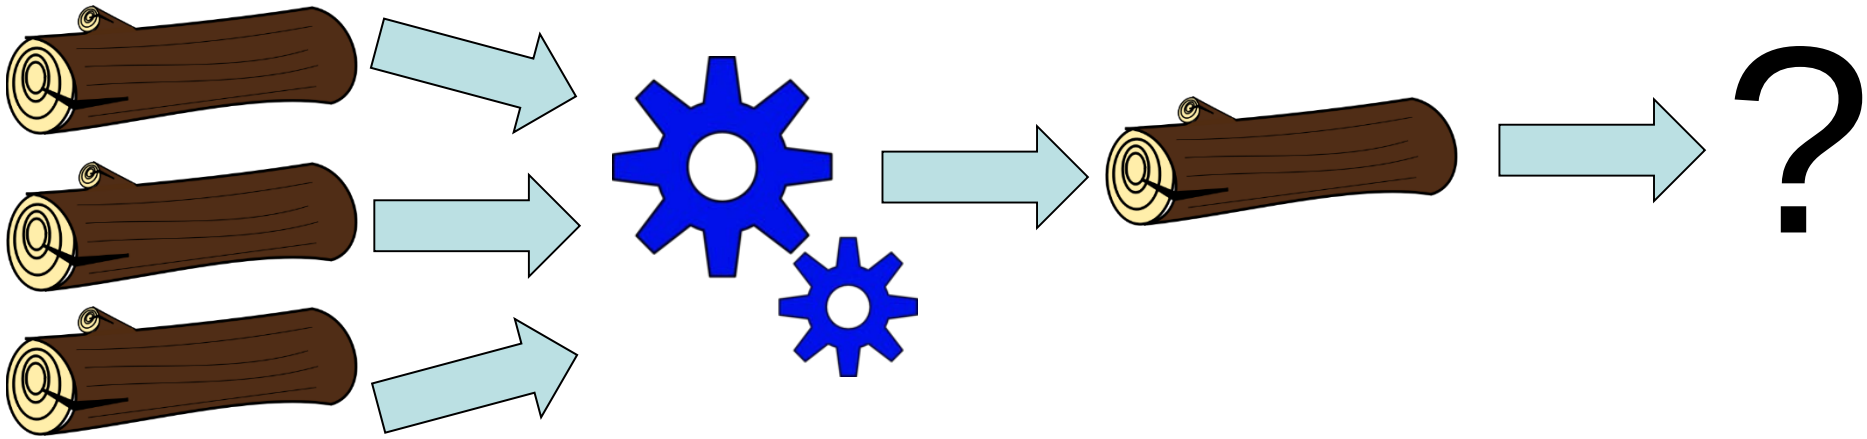

Consequences

- Logs become a persistence mechanism

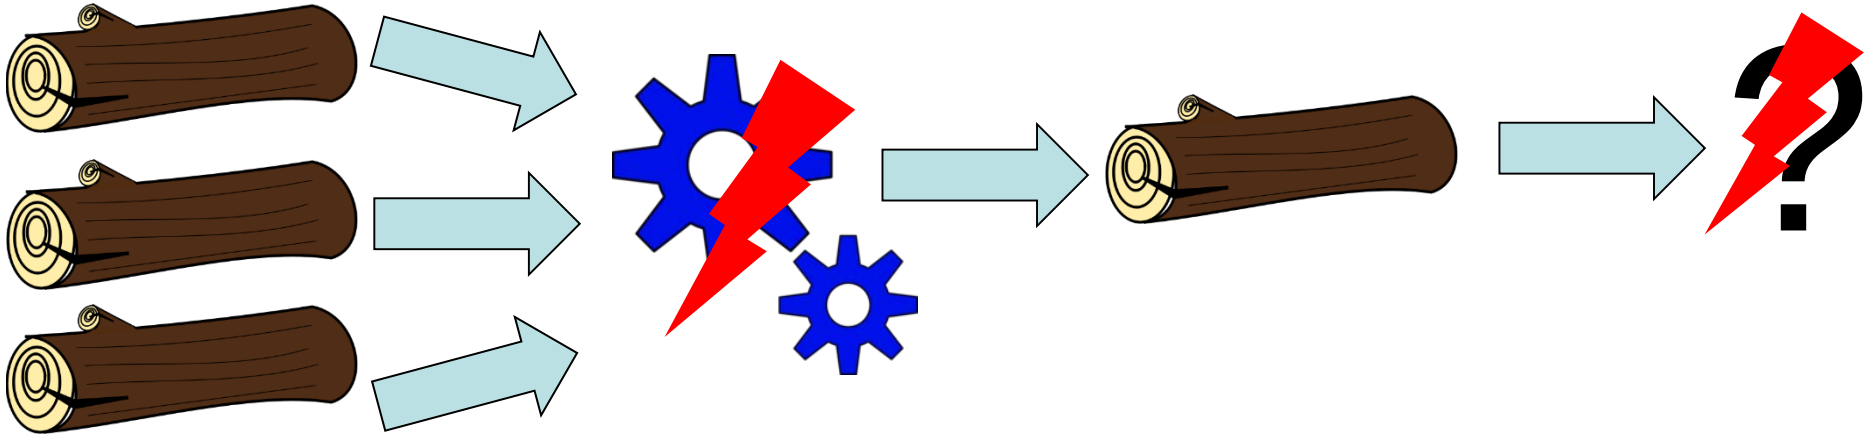

Consequences

- Logs become a recovery / replay mechanism

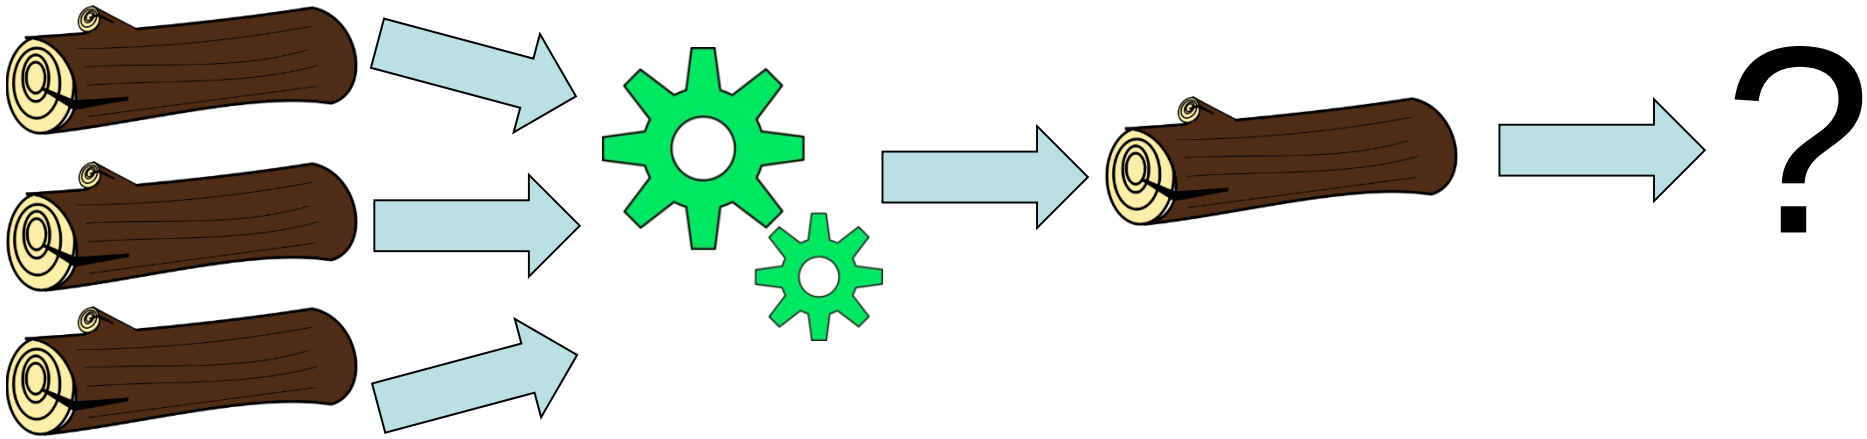

Consequences

- Multiple readers allow “what if”

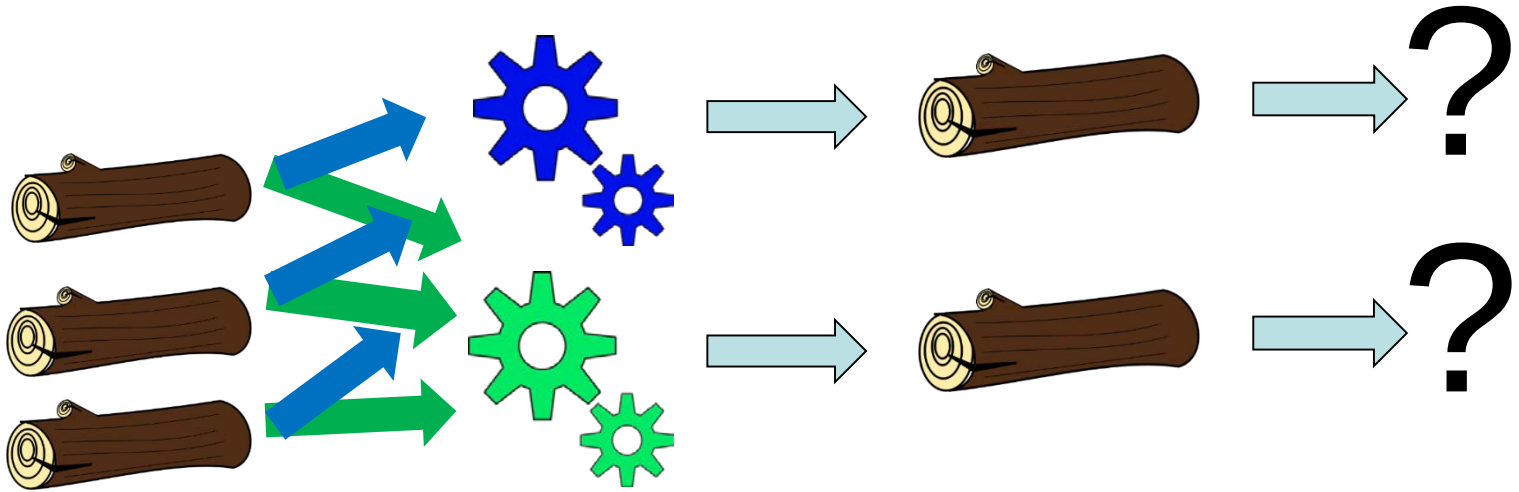

CQRS pattern

- Split out data generation from data querying
- Reads and writes optimized differently
- Can have many read models for one event log

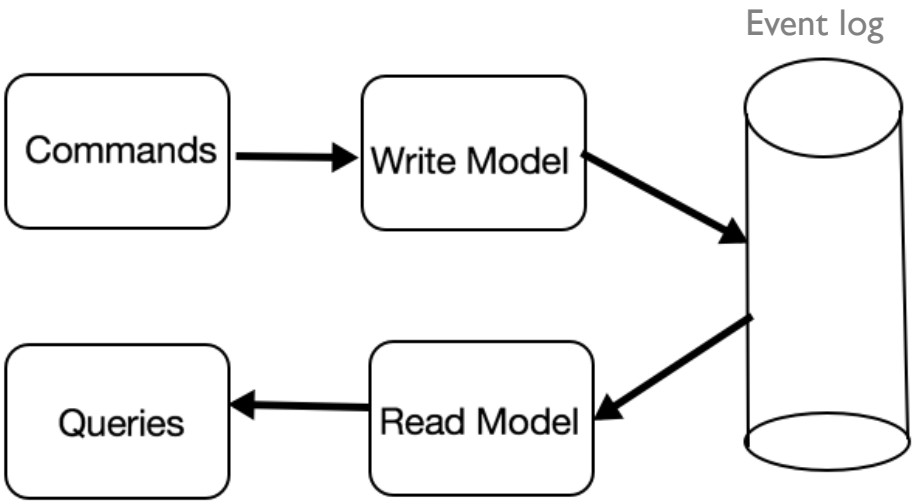

How would this work?

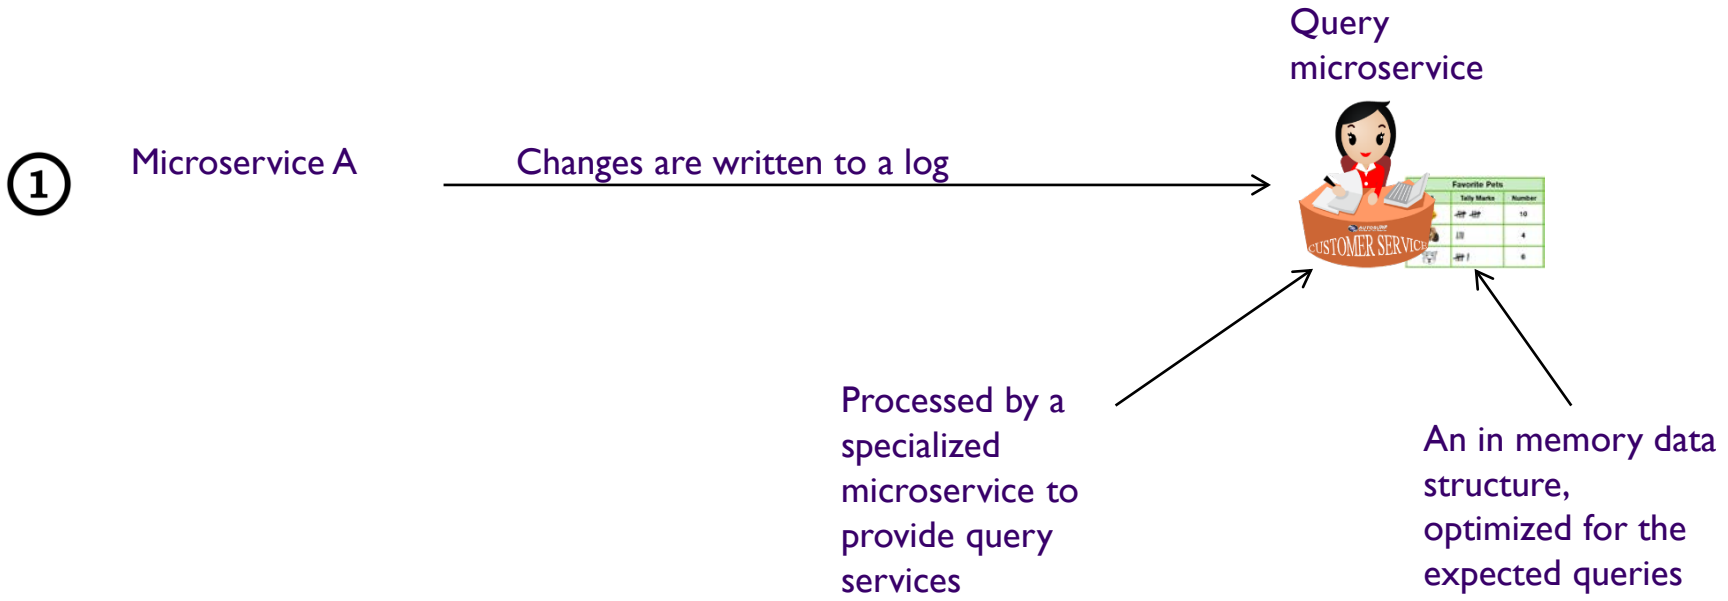

Resiliency is automatic

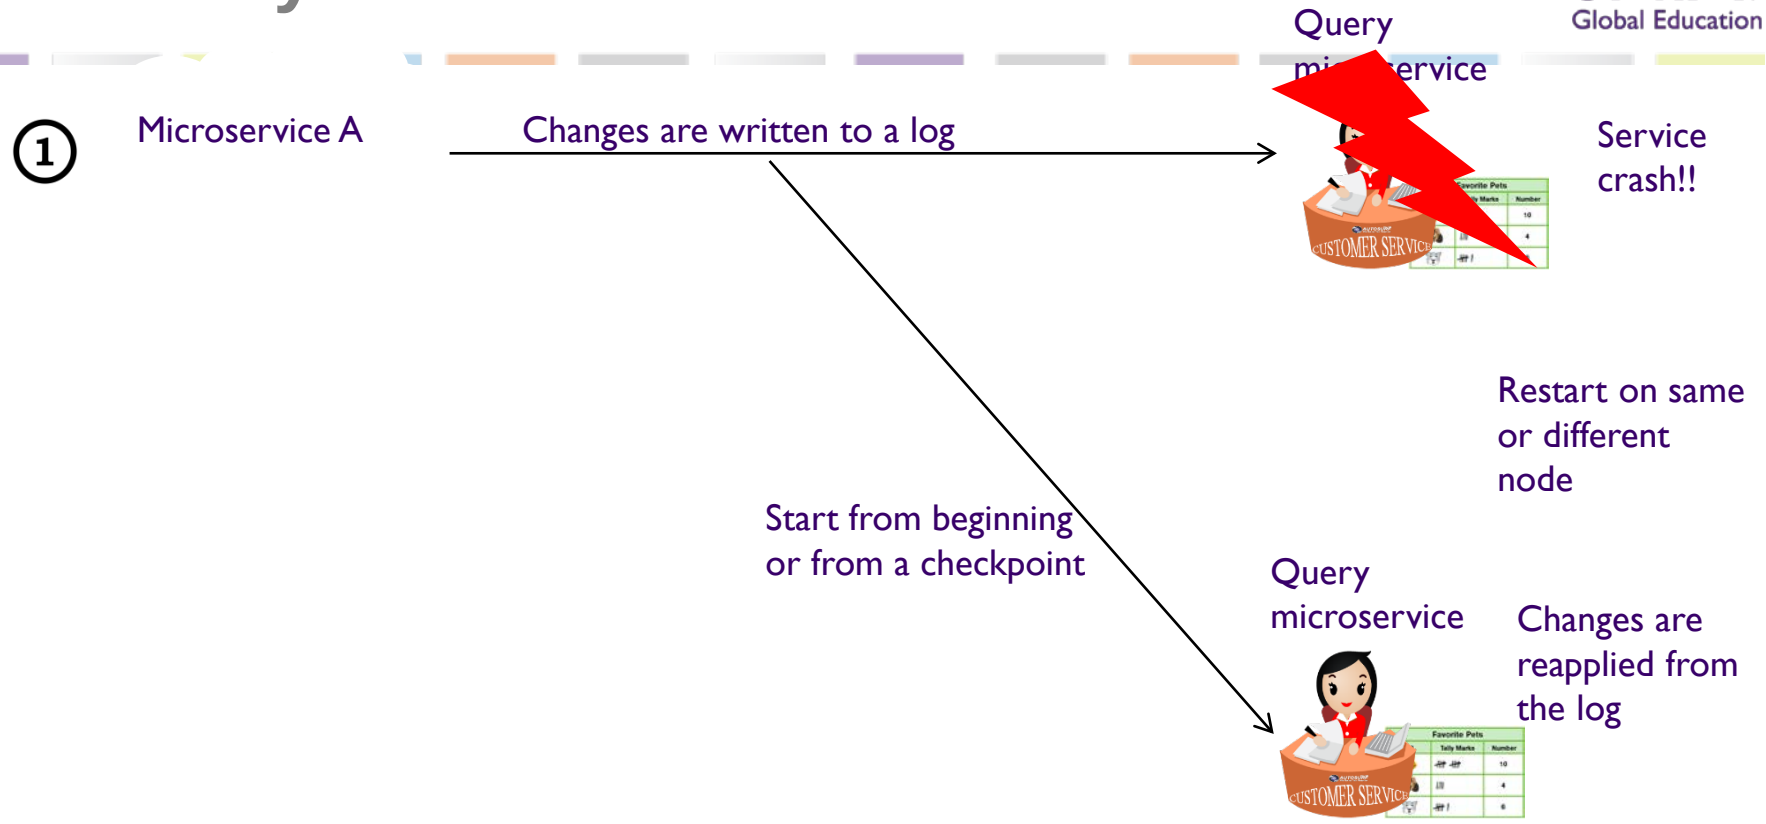

Elastic scaling

① Microservice A → Changes are written to a log

Query microservice

Load grows

Elastic scaling – clone service

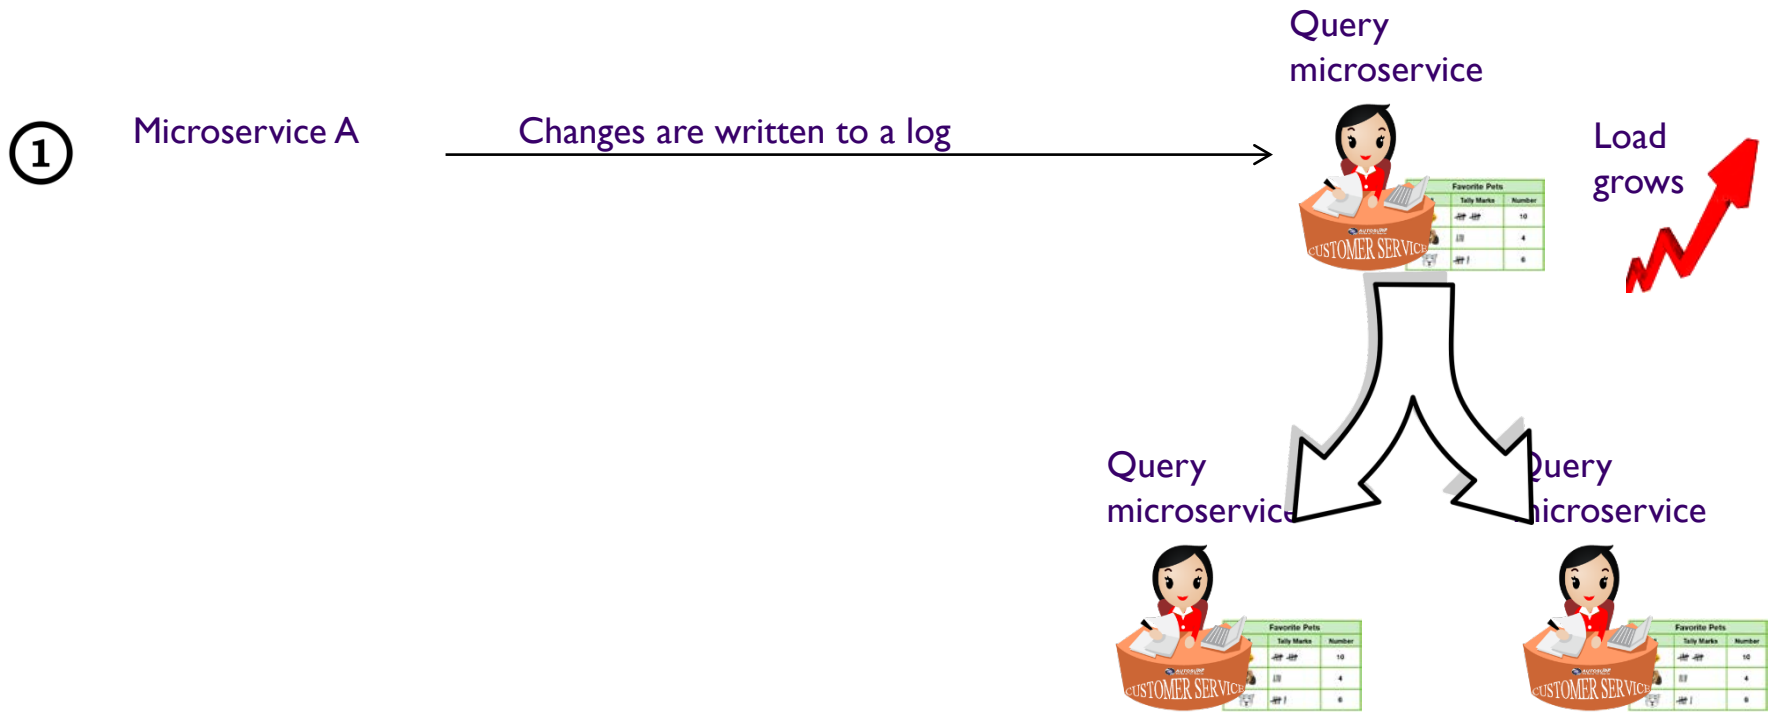

Elastic scaling – partitioned clones

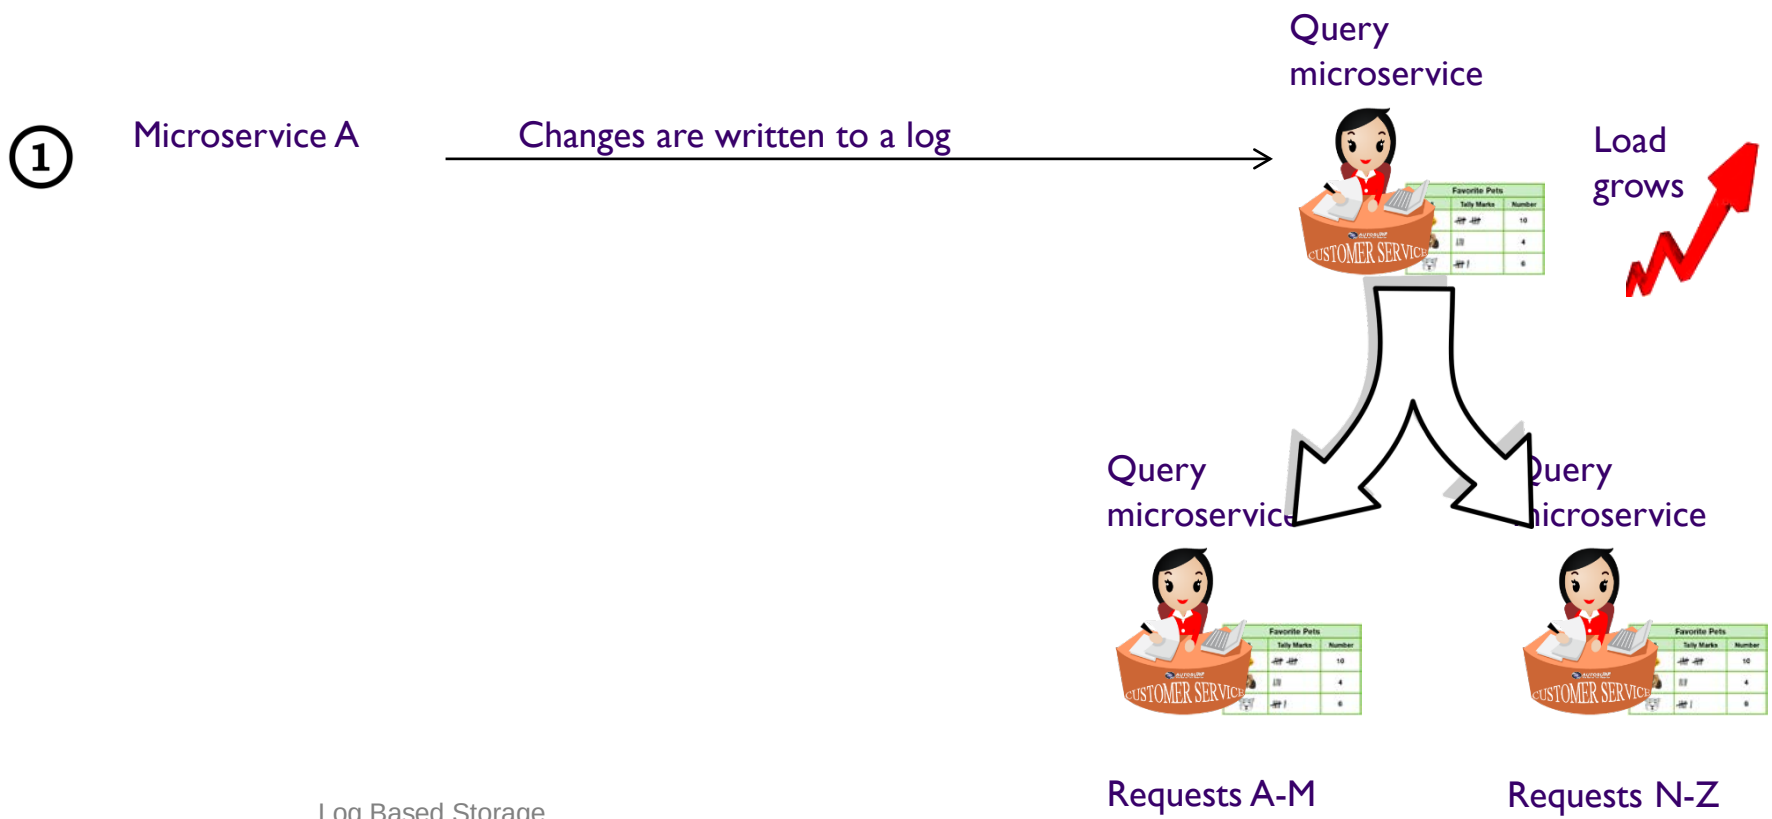

Elastic scaling – clones take over

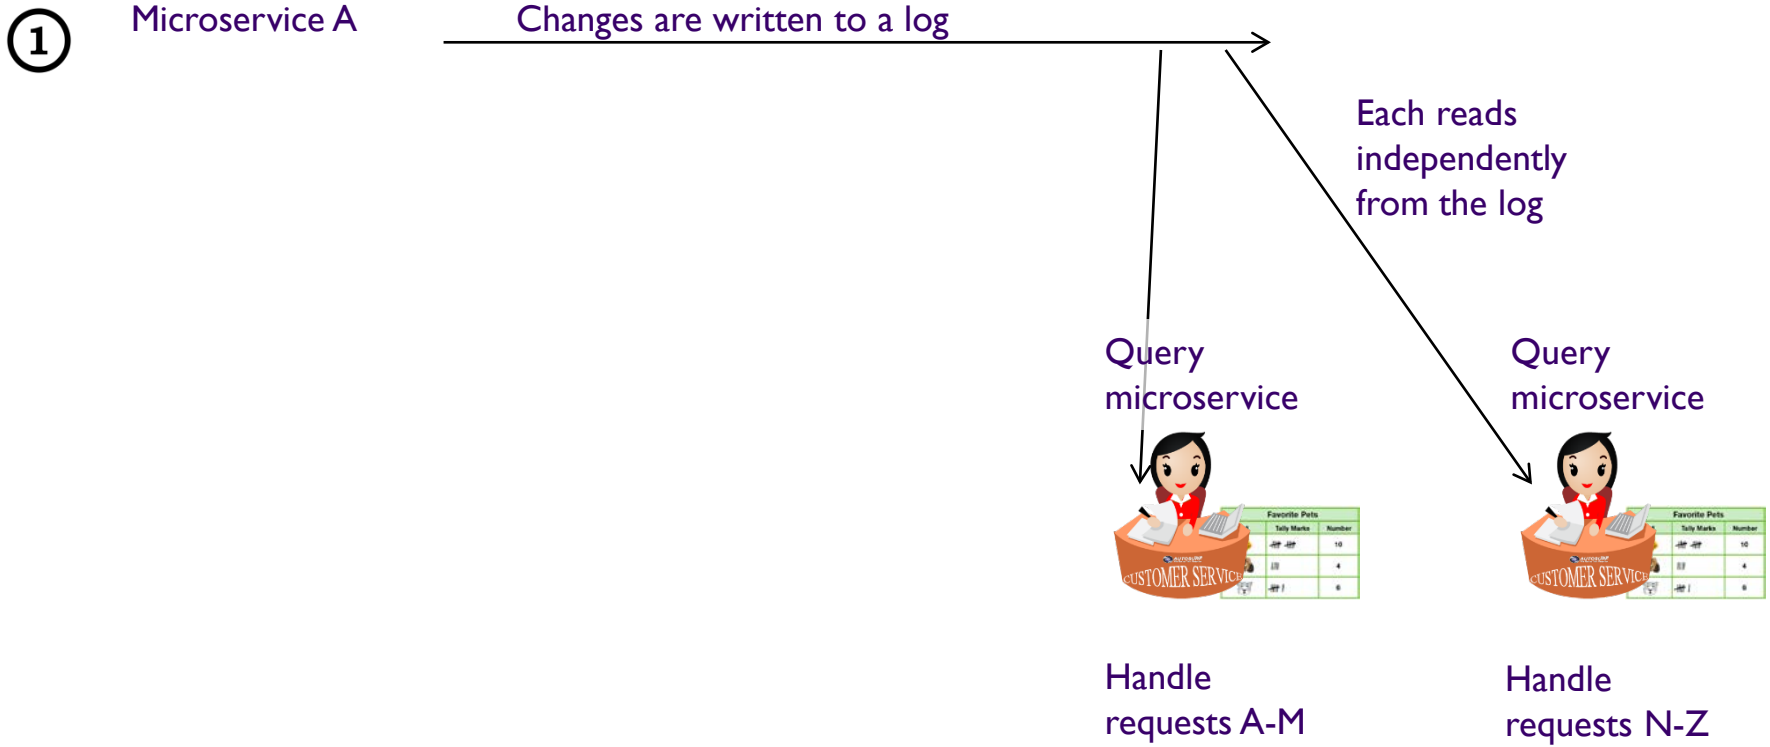

More than just query

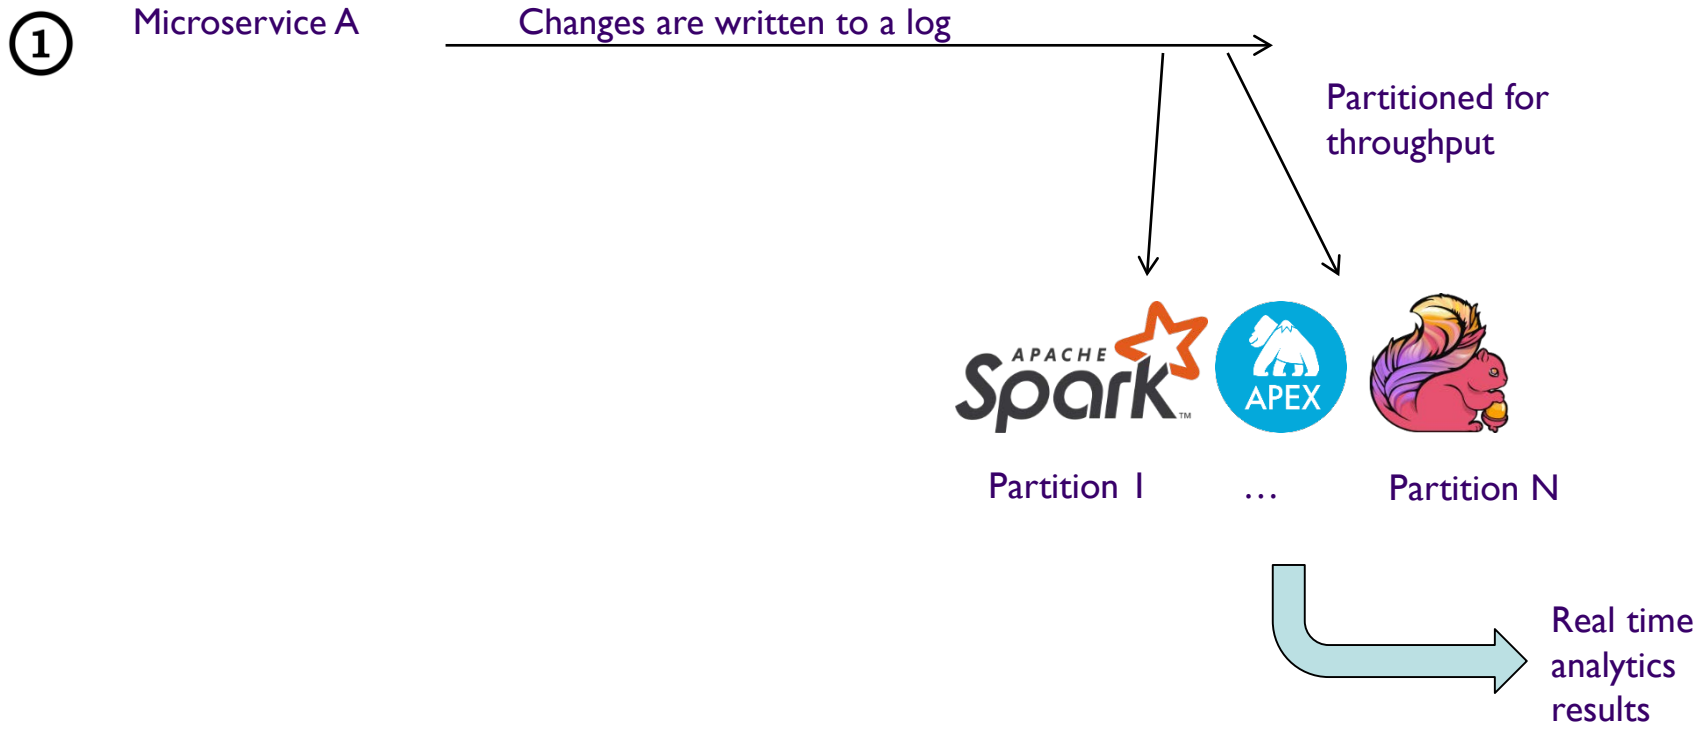

Is log storage truly different?

Block

- Fast - IOPS
- Low overhead
- Application determines data layout
- Inflexible
- Hard to share

File

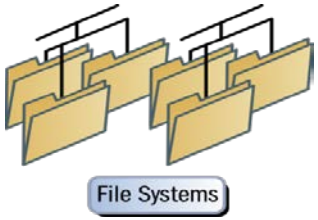

- Human accessible
- Application suggests data layout
- Sharing locally, via NAS
- *Built on block*

Object

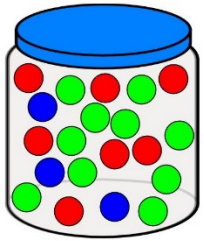

- Human accessible
- Sharing at web scale
- Infinite scale
- BW oriented
- *Built on file normally*

Log

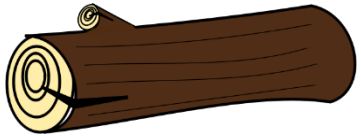

- Human accessible
- Sharing amongst local apps
- Infinite scale
- Recent data: fast
- Older data: BW
- *Build by tiering*

Summary: Logs

- Decoupled from applications
- Source data for multiple applications
- Allow app transparent HA, recovery, replication
- Source data for streaming analytics
- Different characteristics from Block / File / Object
  - Low latency for tail (IOPS) – read and write
  - High bandwidth for others – read only (immutable)

Attribution & Feedback

The SNIA Education Committee thanks the following Individuals for their contributions to this Tutorial.

Authorship History

Mark O'Connell

Updates:

Additional Contributors

Please send any questions or comments regarding this SNIA Tutorial to ***tracktutorials@snia.org***

Tutorial Preparation Guidelines

- Large size slide page number is deliberate
  - Easier for audience to follow from handouts
- Keep slide-count to amount you can cover in time allotted
  - Maximum 40-45 slides for 60-minute session
  - Maximum 30-35 slides for 45-minute session
  - Title, boilerplate (e.g. legal, Attribution & Feedback contents, and 'separator' slides don't count against limit
  - You may have extra slides in an Appendix (max. of ~30% extra)
- 3 levels of bullets maximum
  - 4 to 6 Bullets: use 2 column text slide
  - Use www.snia.org scheme colors for hyperlinks
- AVOID
  - Moving or resizing the text area
  - Changing fonts, point sizes or spacing
- Use color, bold and *italics* for emphasis

Tutorial Preparation Guidelines (cont.)

- Title Slide: List the presenter name/company name; if a co -presenter, list name/company name (no email addresses).
- Always include: Legal, Abstract, and Attribution & Feedback slides (do not delete these slides)
  - Legal slide is #2 slide
  - Abstract Slide is #3: View the sample abstract in "SNIA Tutorial Preparation Guidelines" presentation and modify your abstract to fit that format.
  - Attribution & Feedback (authorship history & contributors) is the final slide before any appendix
- Please try to stay with template colors in case re-coloration is necessary.
  - For information on approved colors, view https://members.snia.org/members/Style%20Guide_07.pdf
- Run spell check when done
  - Choose Tools, Spelling in Powerpoint 97-2003
  - Choose Review, Spelling in Powerpoint 2007
- Save your presentation as a Powerpoint 2007.pptx presentation file
- Refer to the FAQ for SNIA Tutorials "General Information" at http://www.snia.org/education/tutorials/faq/

Graphics

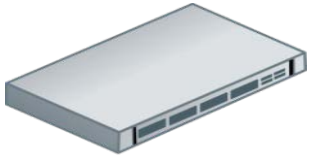

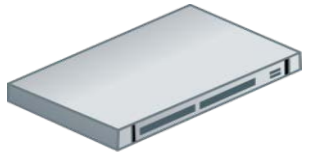

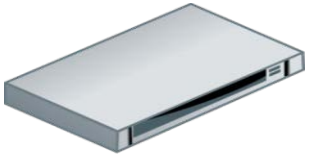

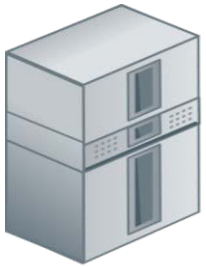

Graphics

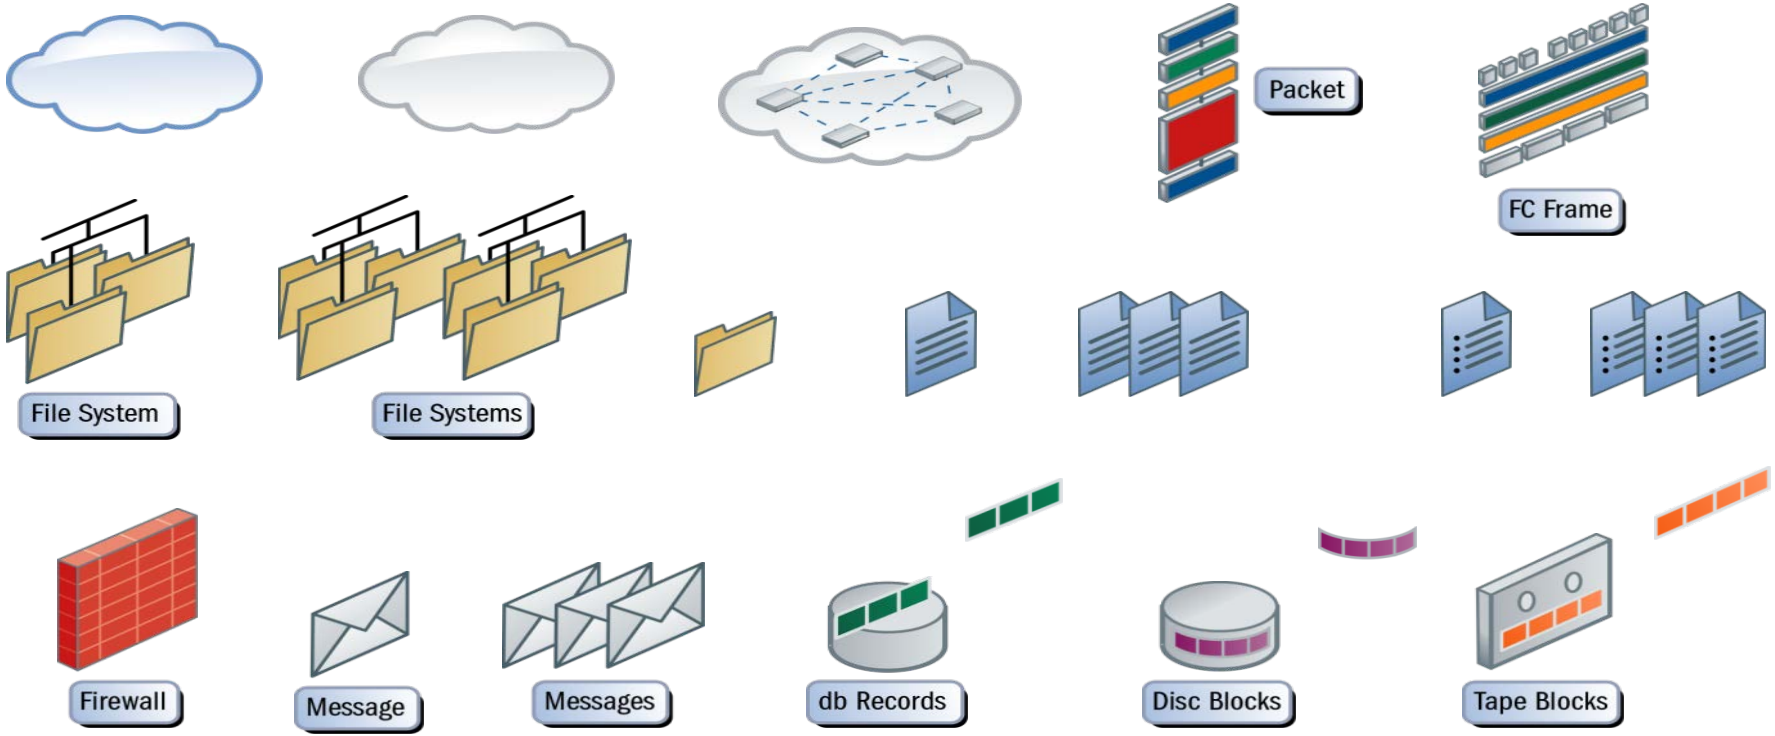

Graphics

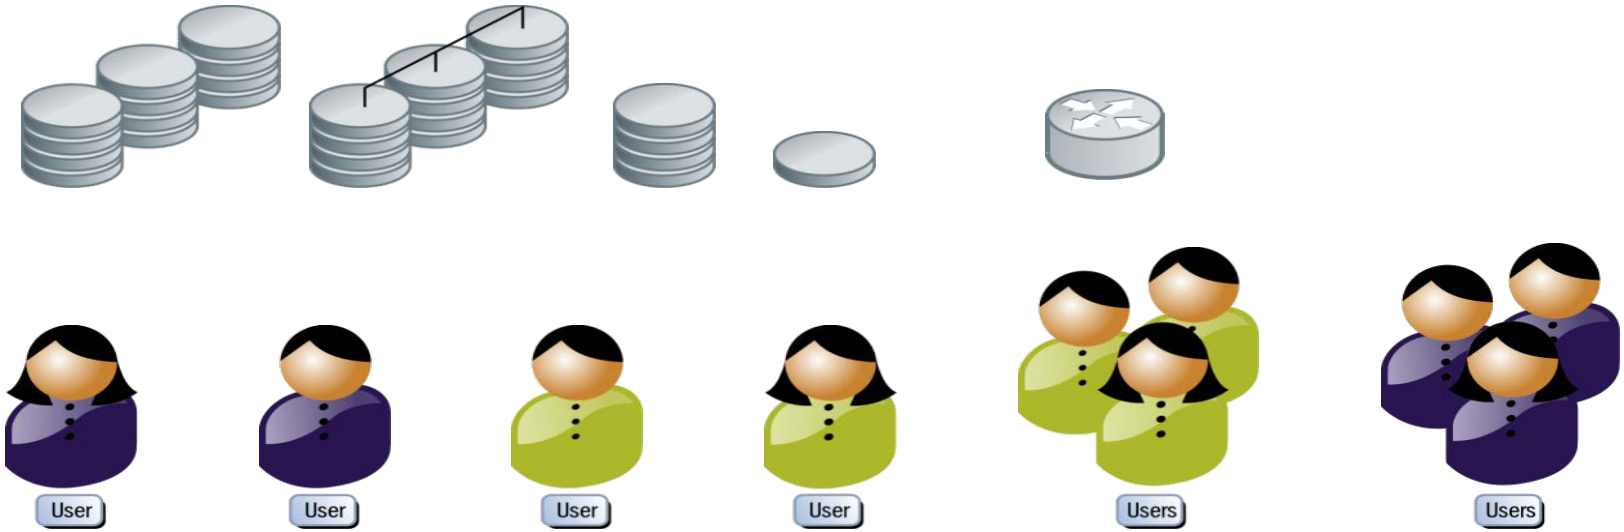

Refer to Other Tutorials

- Use this icon to refer to other SNIA Tutorials where applicable

Check out SNIA Tutorial:

Enter Tutorial Title Here

Attribution & Feedback

The SNIA Education Committee thanks the following Individuals for their contributions to this Tutorial.

Authorship History

Name/Date of Original Author here:

Updates:

Name/Date
Name/Date
Name/Date

Additional Contributors

Name of contributor here
Name of contributor here
Name of contributor here
Name of contributor here
Name of contributor here
Name of contributor here
Name of contributor here
Name of contributor here

Please send any questions or comments regarding this SNIA Tutorial to ***tracktutorials@snia.org***