

Welcome to the IBTA Fall Webinar Series!

A four-part webinar series devoted to making I/O work for you

Why I/O is Worth a Fresh Look

Presented by the InfiniBand Trade Association

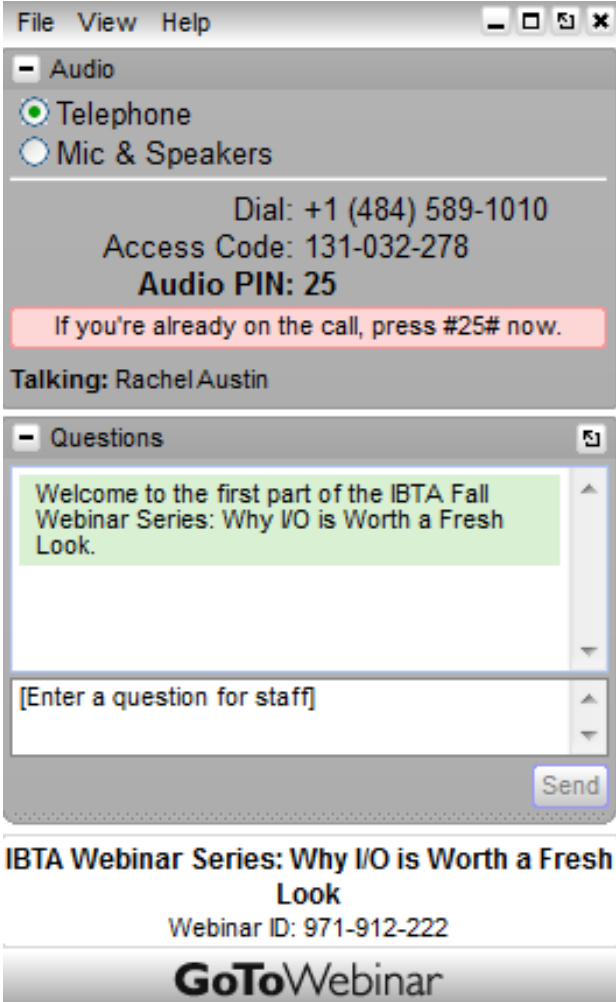


IBTA Fall Webinar Series

- ➔ September 23 Why I/O is Worth a Fresh Look
- October 21 The Practical Approach to Applying InfiniBand in Your Data Center
- November 11 InfiniBand Technology: No Magic, Just Solid Engineering
- December 9 An Expanding Role for RDMA in Future Data Centers

Webinar Logistics

- All attendees are muted
- Listen via your computer speakers or telephone
 - Audio broadcast through your computer speakers is the default
 - To listen by telephone, dial the phone number in your invitation or the number displayed in your control panel
- Submit questions via the Questions pane in your control panel
 - Questions will be addressed at the end of the webinar
- A recording of the webinar will be available at www.infinibandta.org



The screenshot shows a software interface for a webinar. At the top is a menu bar with 'File', 'View', and 'Help'. Below it is an 'Audio' section with two radio buttons: 'Telephone' (selected) and 'Mic & Speakers'. The telephone section displays the dial number '+1 (484) 589-1010', the access code '131-032-278', and the audio PIN '25'. A red highlighted box contains the text 'If you're already on the call, press #25# now.' Below this, it says 'Talking: Rachel Austin'. The 'Questions' section has a green message box that reads 'Welcome to the first part of the IBTA Fall Webinar Series: Why I/O is Worth a Fresh Look.' There is a text input field with the placeholder '[Enter a question for staff]' and a 'Send' button. At the bottom, the webinar title 'IBTA Webinar Series: Why I/O is Worth a Fresh Look' and 'Webinar ID: 971-912-222' are displayed, along with the 'GoToWebinar' logo.



Paul Grun

Chief Scientist, System Fabric Works
pgrun@systemfabricworks.com



Jim Ryan

Intel
jim.ryan@intel.com



We need better results, more quickly...and it can't cost more.

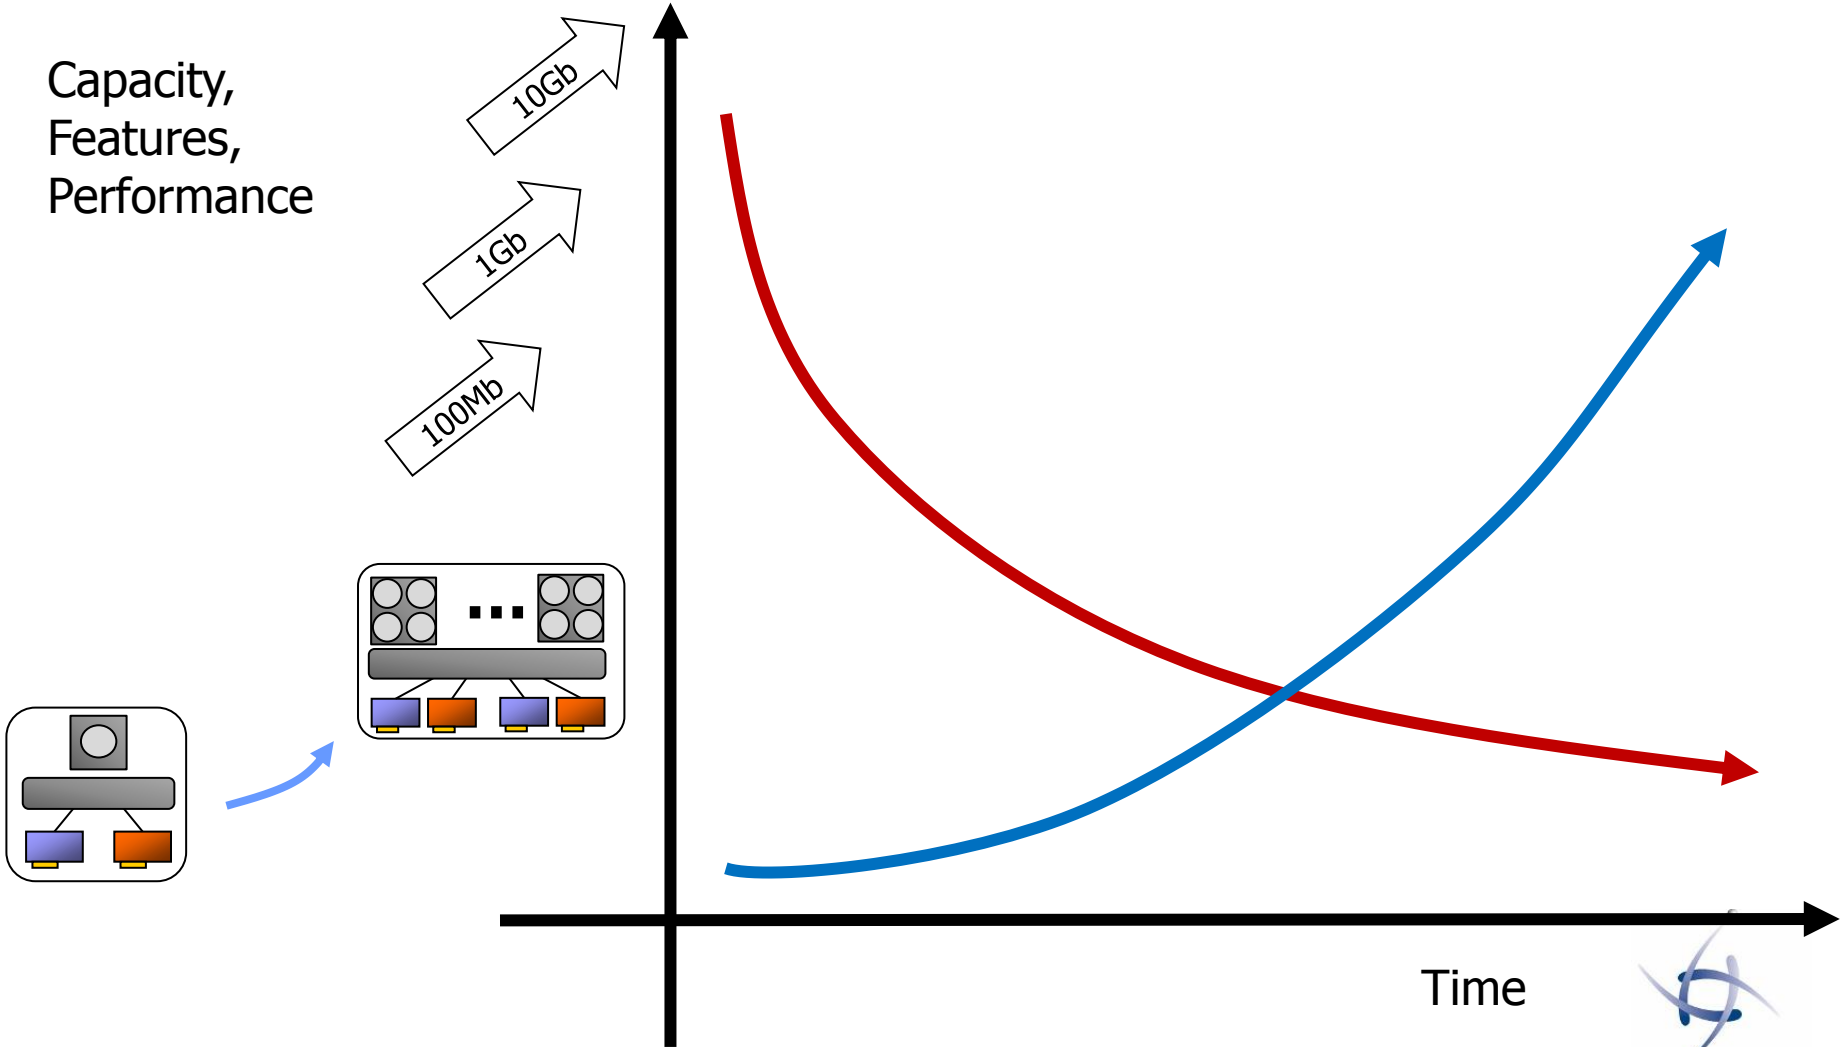


Between A Rock and A Hard Place

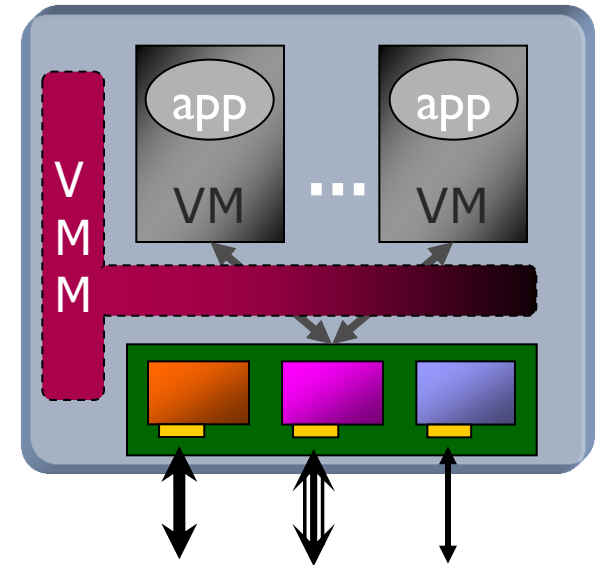
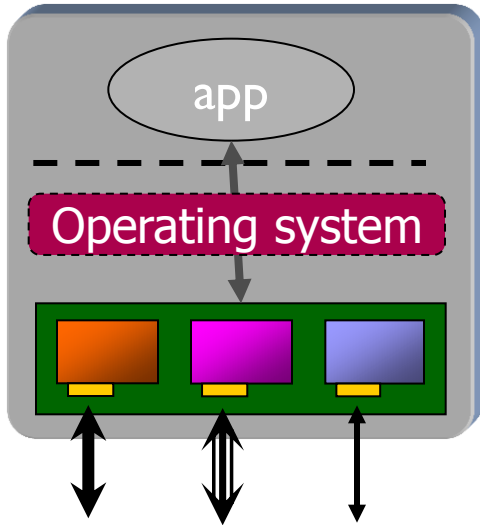


- Keep your users happy
- With less resources too
- Technology to the rescue

The Technology Curve


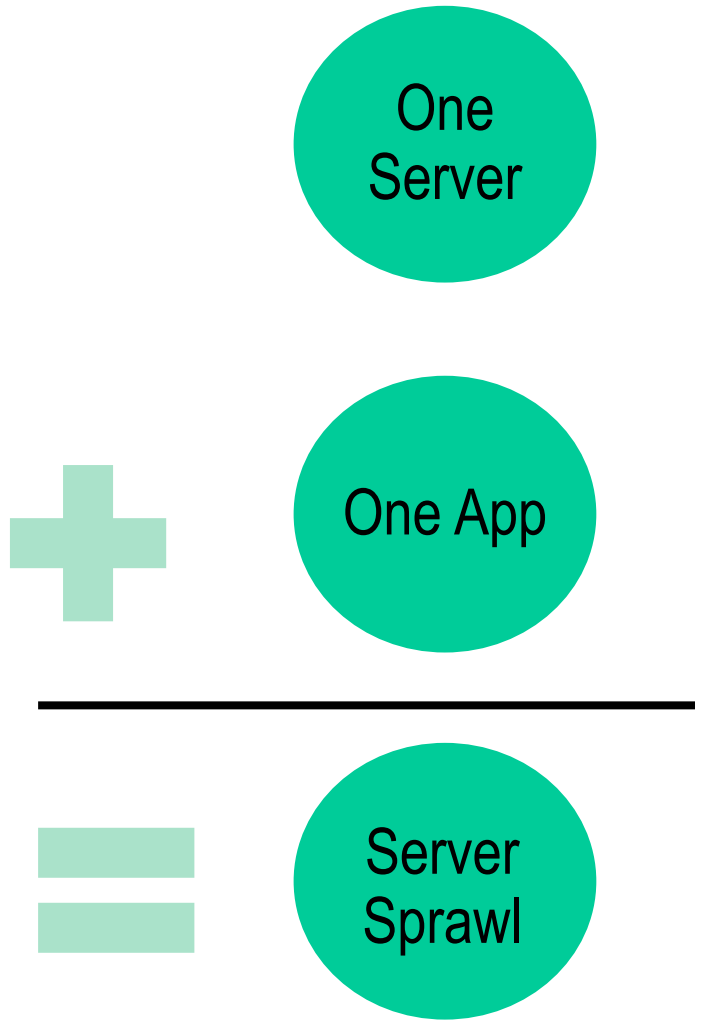


Virtualization

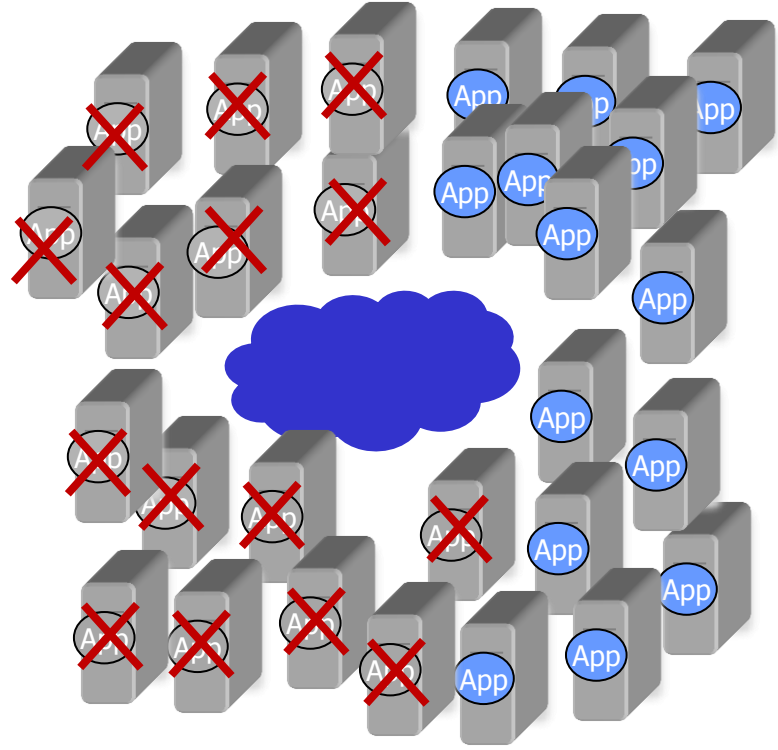



Virtualization allows fine-grained allocation of server resources
→ improve server utilization

Server Sprawl



Server Sprawl




I/O – Subtle Changes



- Migration from parallel buses to serial buses
- A pretty disruptive change
- But server platforms are more performant than ever

I/O – Subtle Changes

Major speed bumps in interconnect speeds



Is It Still the Right Model for I/O?



- Do we have the right model for an I/O interface?
- The kernel as a component of the I/O subsystem?
- Data copies?

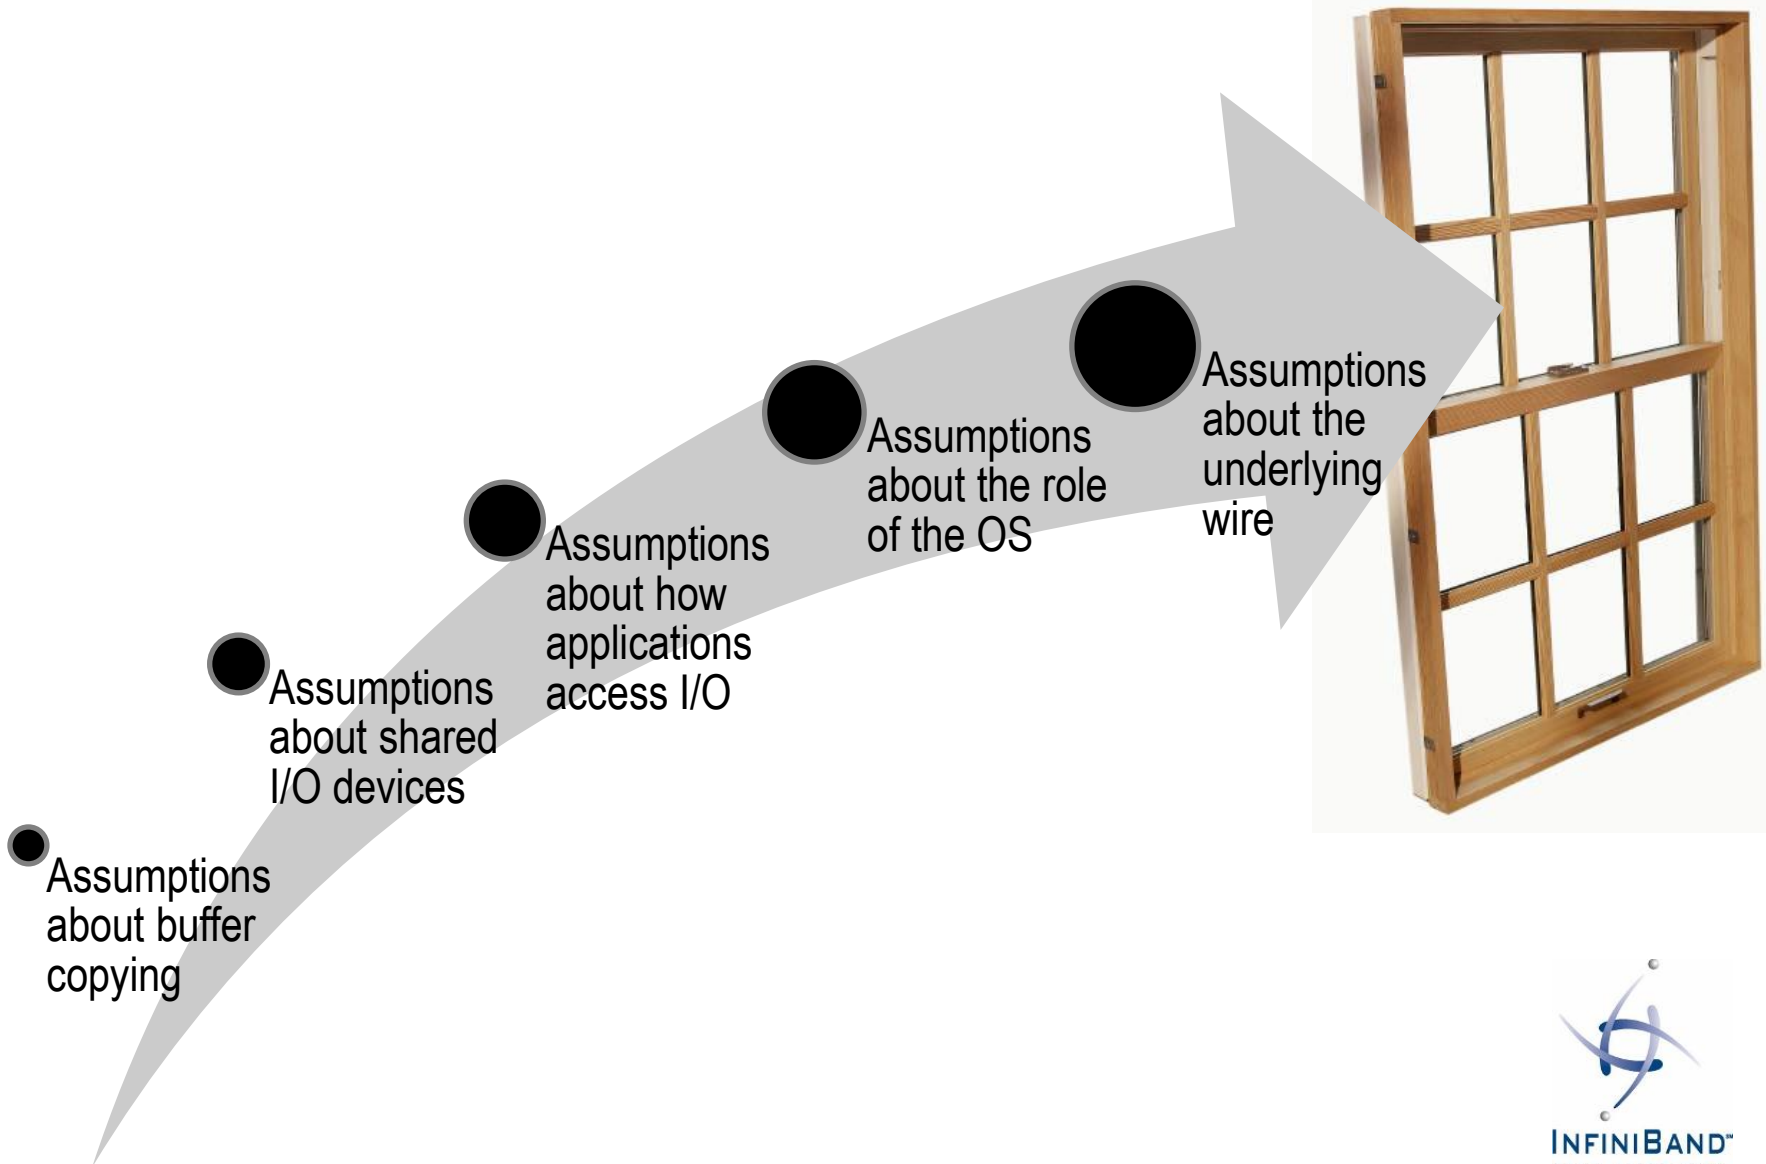


Figures of Merit

1. Original motivation: Solve significant problems in clustering (IPC)
 - Figure of merit: latency
2. Faster networks = faster packet rate = less protocol processing time
 - Trouble ahead for software-based network protocols
 - Figure of merit: CPU utilization
3. Storage: We need to move data fast, and lots of it
 - Figure of merit: bandwidth



How To Innovate in One Easy Step

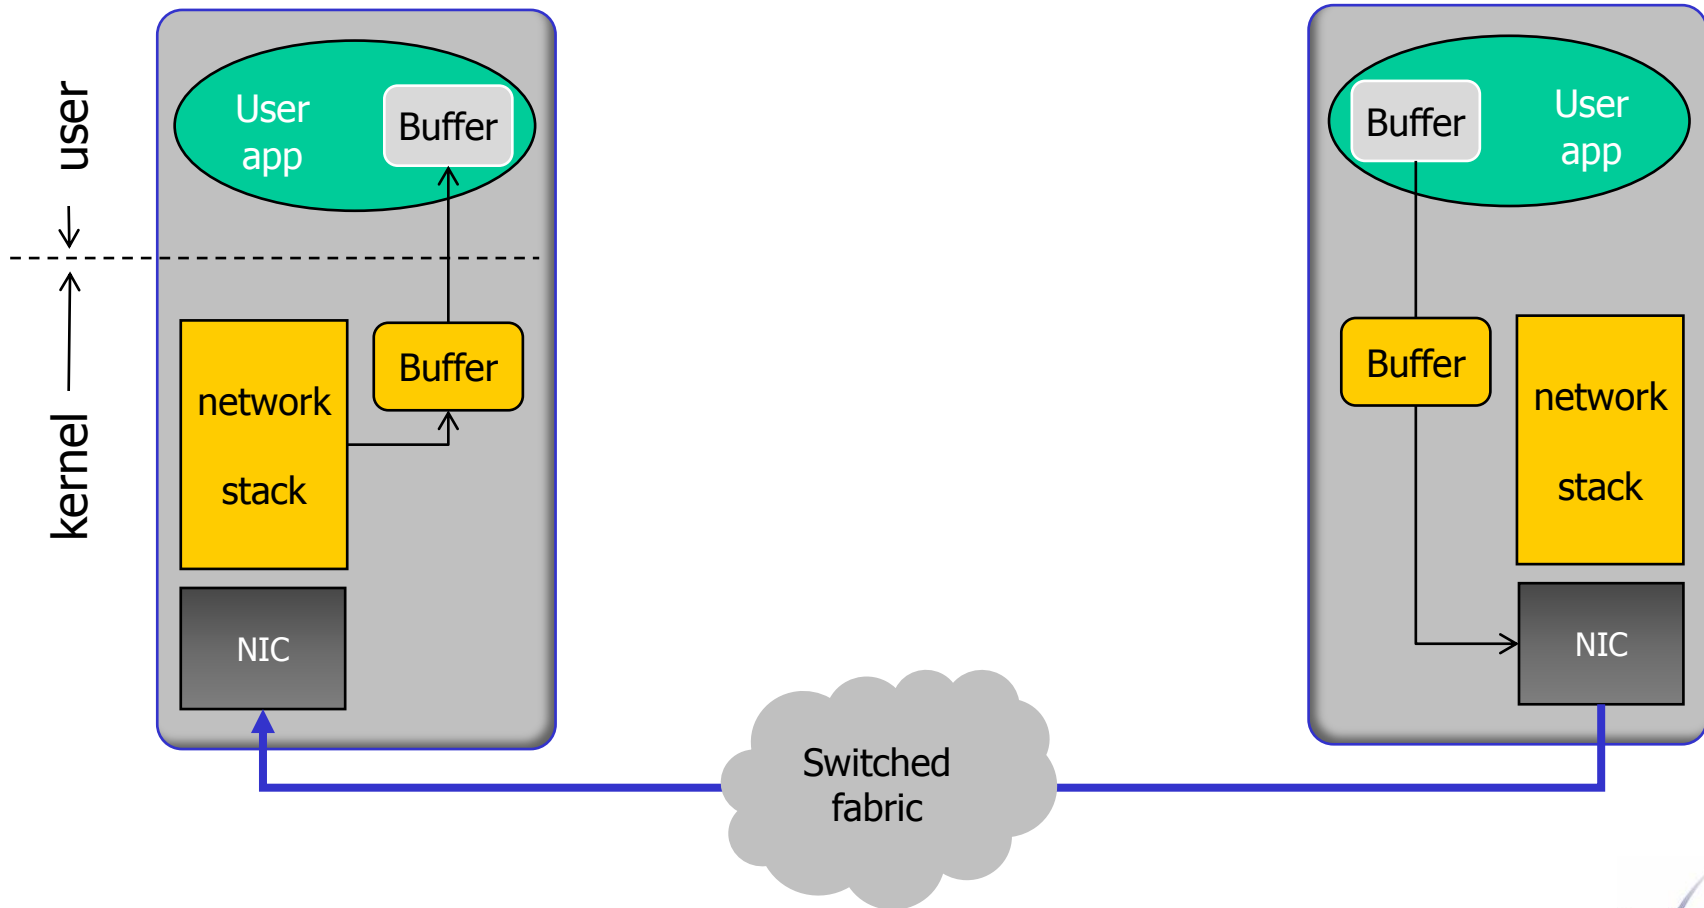


- RDMA – Remote Direct Memory Access

- Applications are given *DIRECT* access to a message passing service
- No need to involve OS
- All I/O operations are conducted as messages




Familiar Network Architecture



■ hardware ■ OS ■ user software



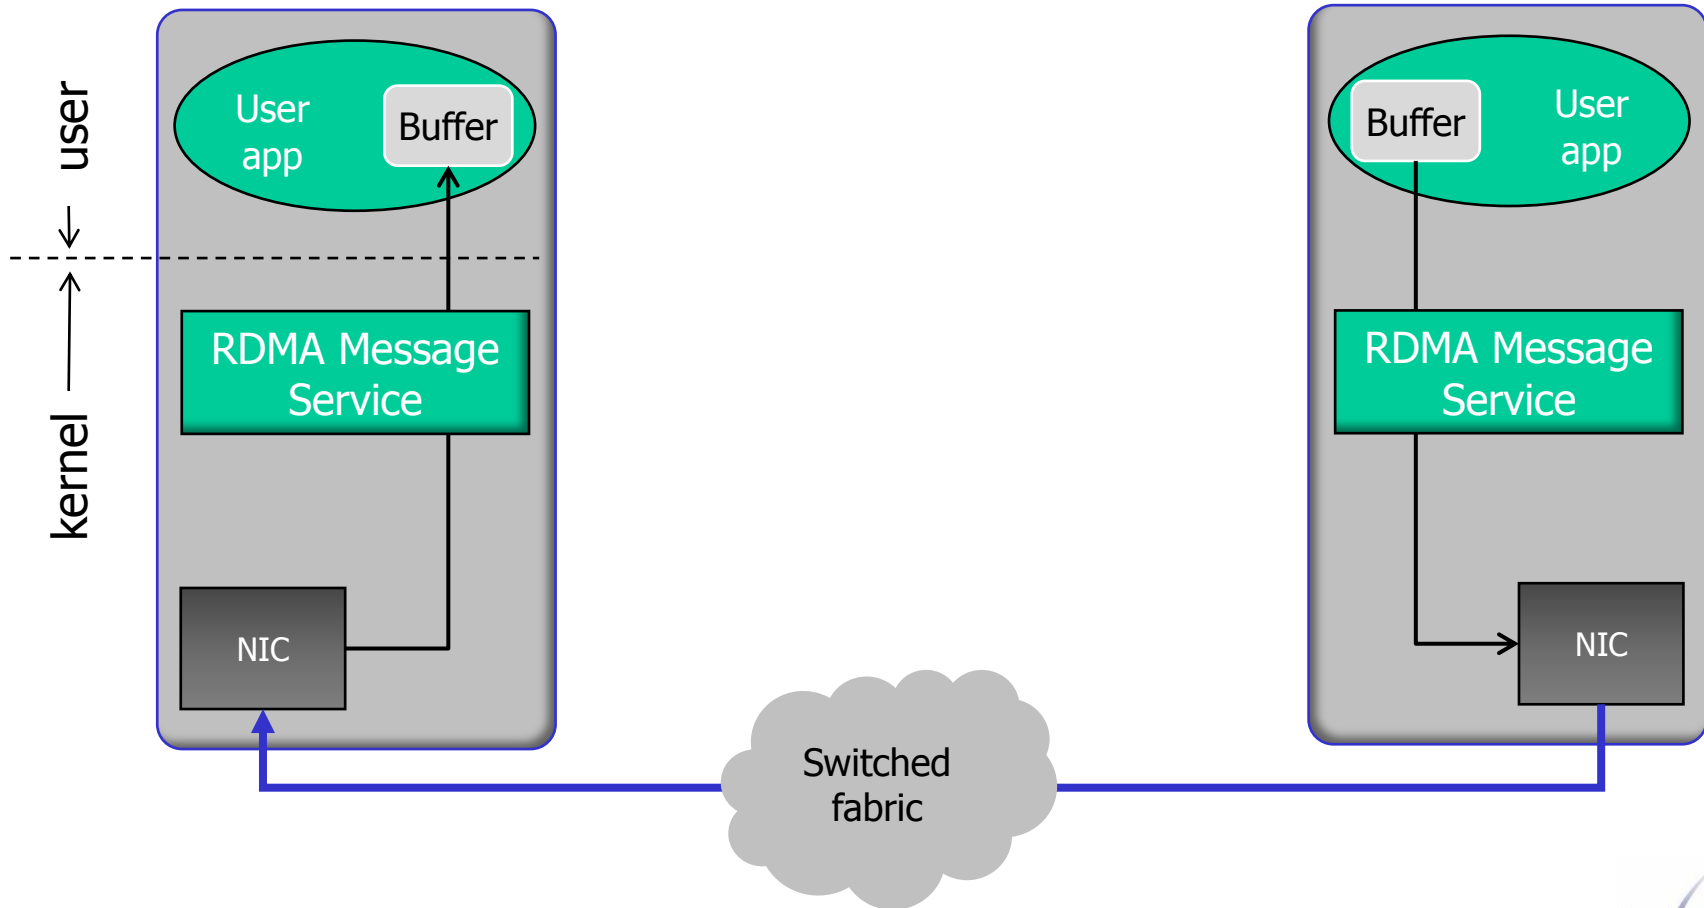
RDMA Network Architecture



■ hardware ■ OS ■ user software




RDMA Network Architecture



■ hardware ■ OS ■ user software



RDMA Network Architecture




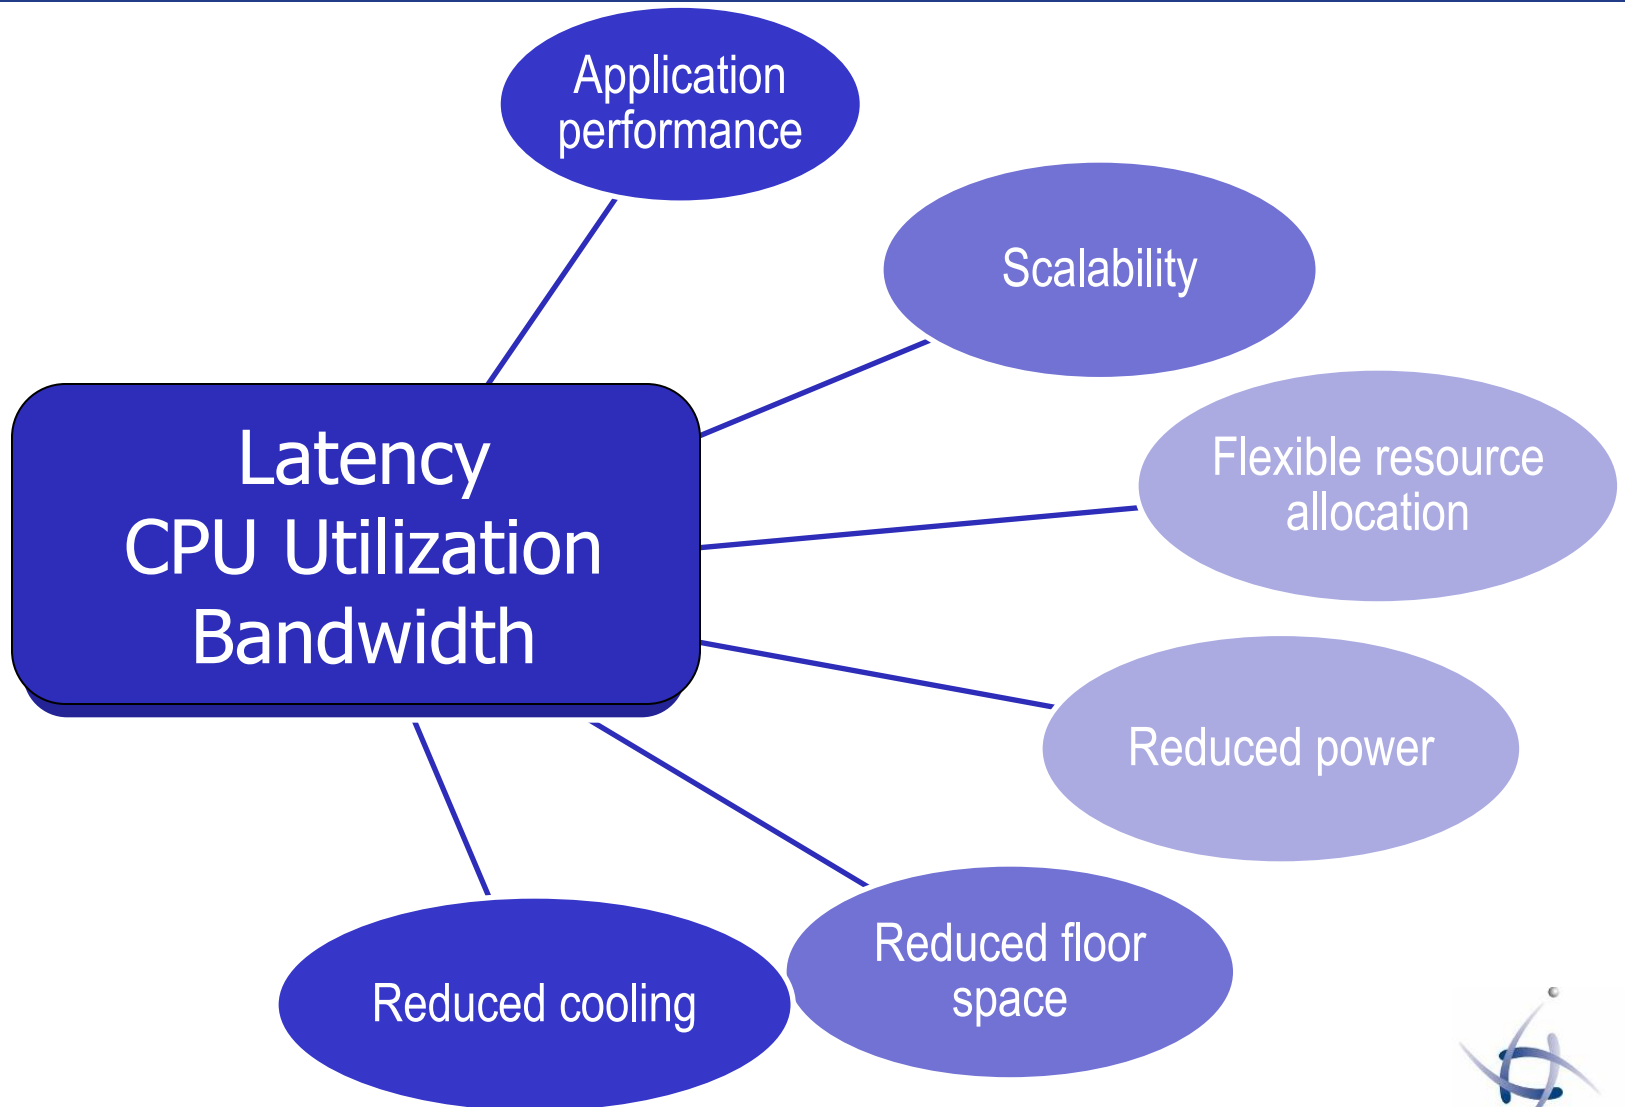
■ hardware ■ OS ■ user software



Efficient, Application-Level I/O


- No context switch for data transfers
- Direct application access to the RDMA message service
- No kernel buffer copies
- Application accesses the NIC directly
- Application directly addresses remote application
- Automatic virtual-to-physical address translation






Evolving the Data Center

Traditional data center network design: servers interconnected by networks

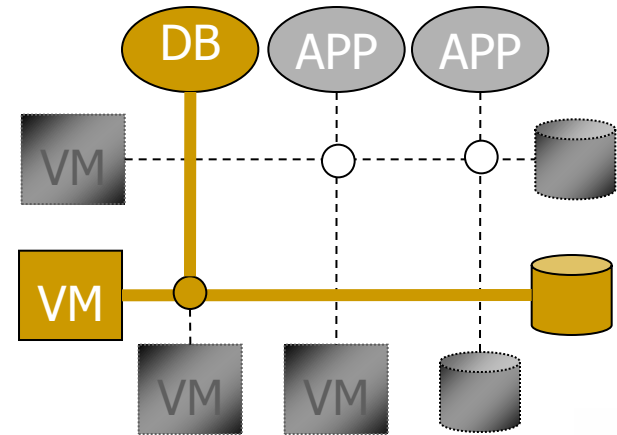


Evolving the Data Center

Traditional data center network design: servers interconnected by networks



Grid-based data center design



Interconnect takes on a new role



Next Webinar: October 21

- The Practical Approach to Applying InfiniBand in Your Data Center
 - October 21
 - 11 am ET/ 8am PT
- RDMA is a disruptive technology
 - So was PCI-e, virtualization, etc.
- This webinar explores some ways to begin the process of integrating an important I/O technology into your environment



Questions?

If we do not answer your question today, please email pgrun@systemfabricworks.com.

For more information about the InfiniBand Trade Association, please visit www.infinibandta.org, and for information about becoming a member, please email administration@infinibandta.org.

