

A First Look at Problems in the Cloud

Theophilus Benson*, Sambit Sahu⁺,
Aditya Akella*, Anees Shaikh⁺

*University of Wisconsin, Madison

⁺IBM Research

Cloud Computing

- New delivery and consumption model for IT
 - No set up cost for users and on-demand usage
 - Several service models emerging:
Infrastructure, Platform, Software.
- Infrastructure-as-a-Service
 - Easily acquire and release virtual servers, storage, bandwidth
 - Provider manages the virtual infrastructure resources
 - User manages the application and OS
- New model = new challenges for service management
 - Virtualized and abstracted resources
 - Limited visibility into the infrastructure layer



What is required to support problem determination in the cloud environment

Problems in the Cloud


- **Our objective**: develop an understanding of the nature of problems experienced by customers of an IaaS cloud

- Understand the types of problems arising in deployment and operation
 - How do these change over time (longitudinal study)
- How users go about solving their problems today
 - Which problems require more help from cloud providers
 - Which problems are persistently difficult to solve
- Develop a preliminary characterization of the forum-based cloud support model
 - See paper for details

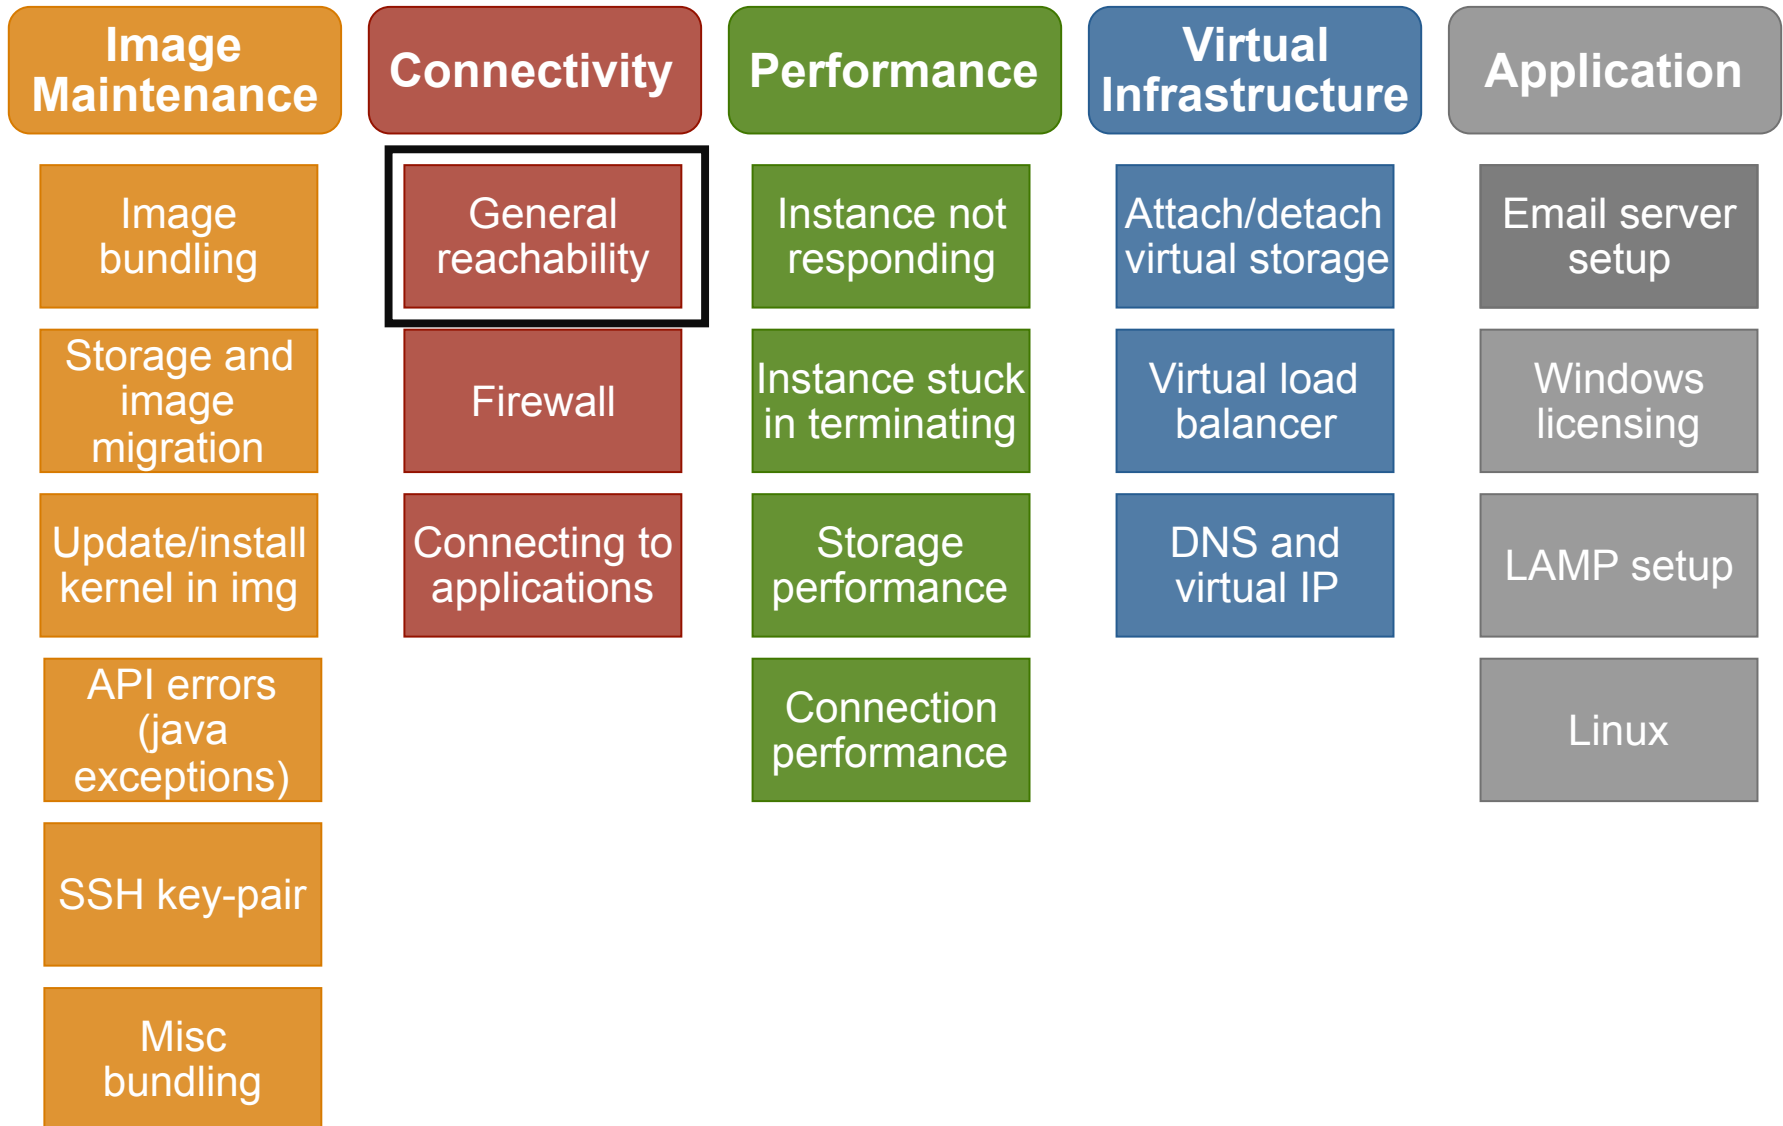
Identify useful mechanisms and best practices for efficient problem resolution in the cloud

Data-driven Approach

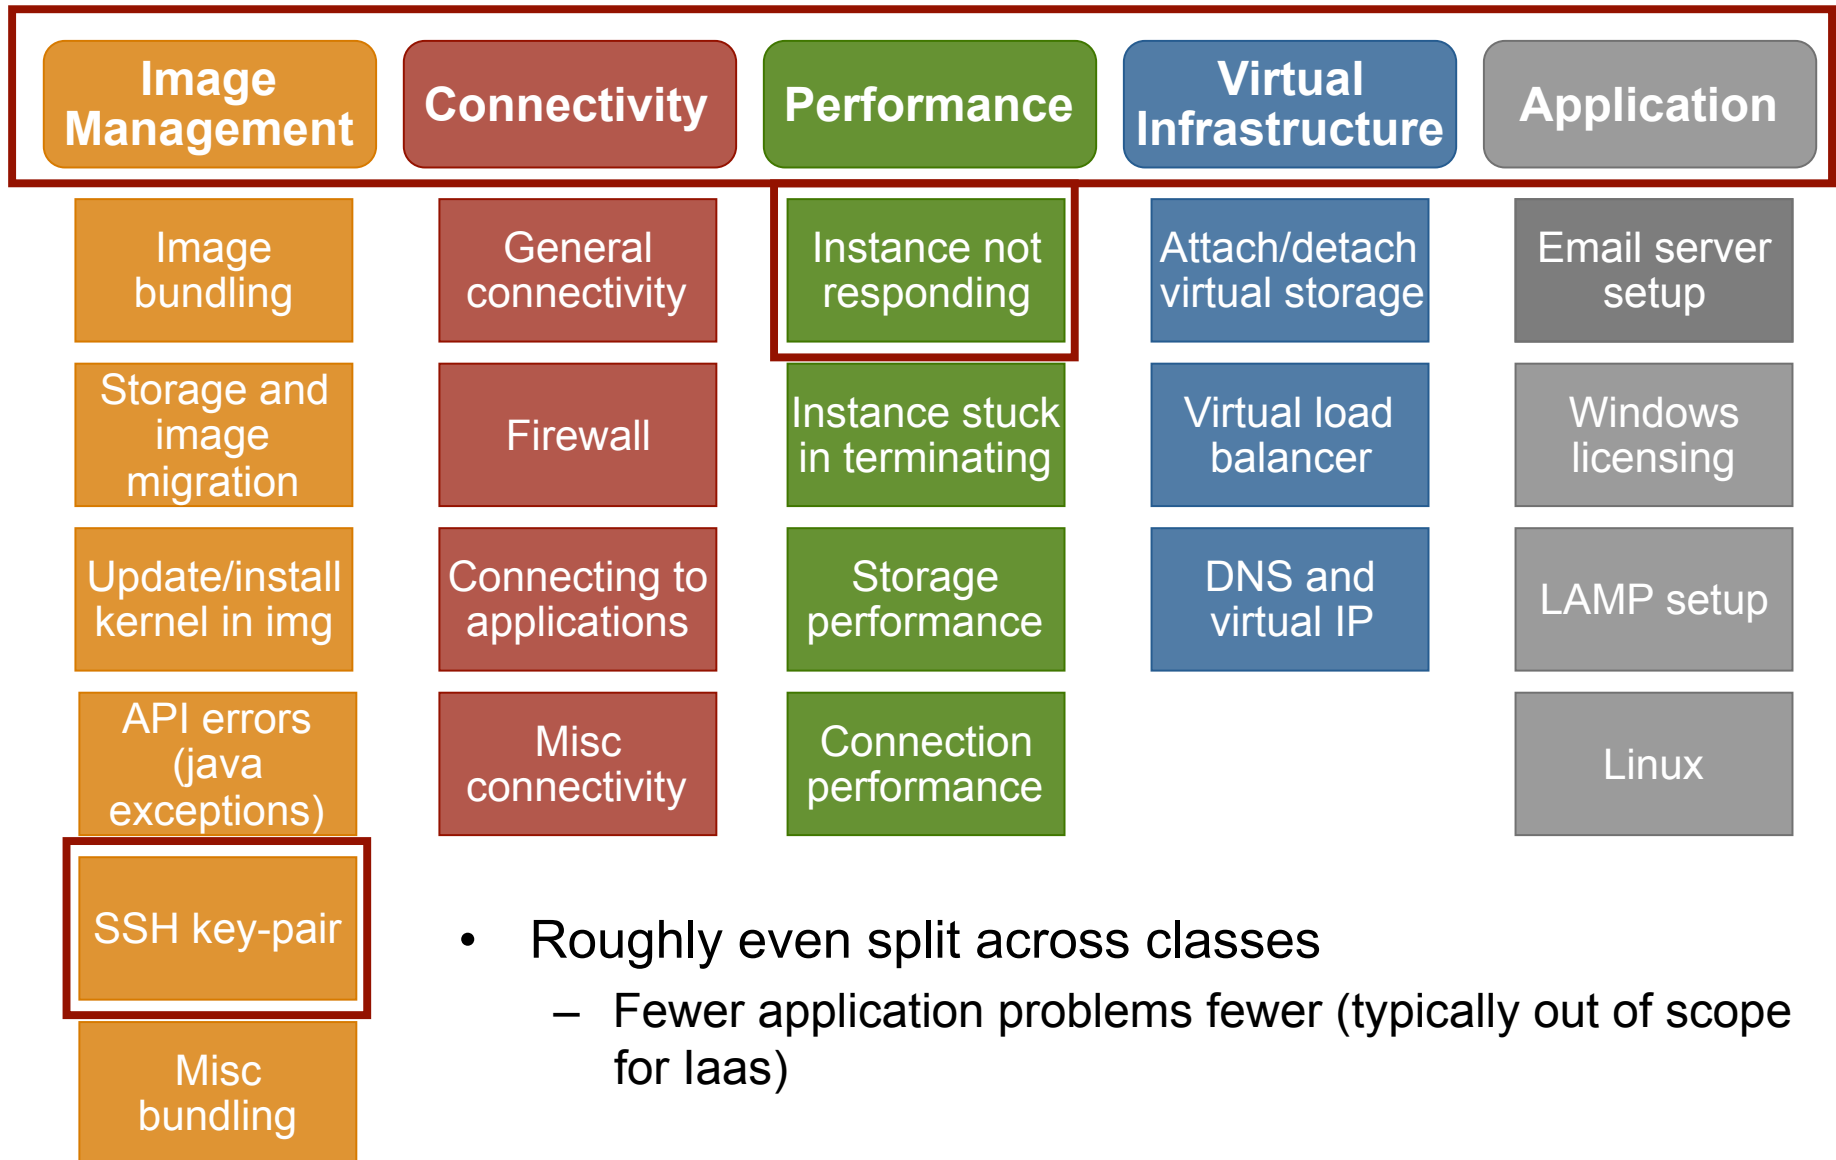
- Study actual user problems and experiences
 - Based on open support forum of a large IaaS Cloud provider
- 3 years of message threads
 - August 2006 through December 2009
 - 9,575 reported problems (threads)
 - Too many for manual classification
 - Each thread consists of a number of related forum postings
 - Free text problem symptoms and suggested solutions
 - Timestamps and user id
 - Need to impose structure for automated grouping
- Text analytics to classify reported problems
 - IR techniques to group problems into clusters
 - 91% of threads mapped to 27 clusters
 - Grouped in 5 classes



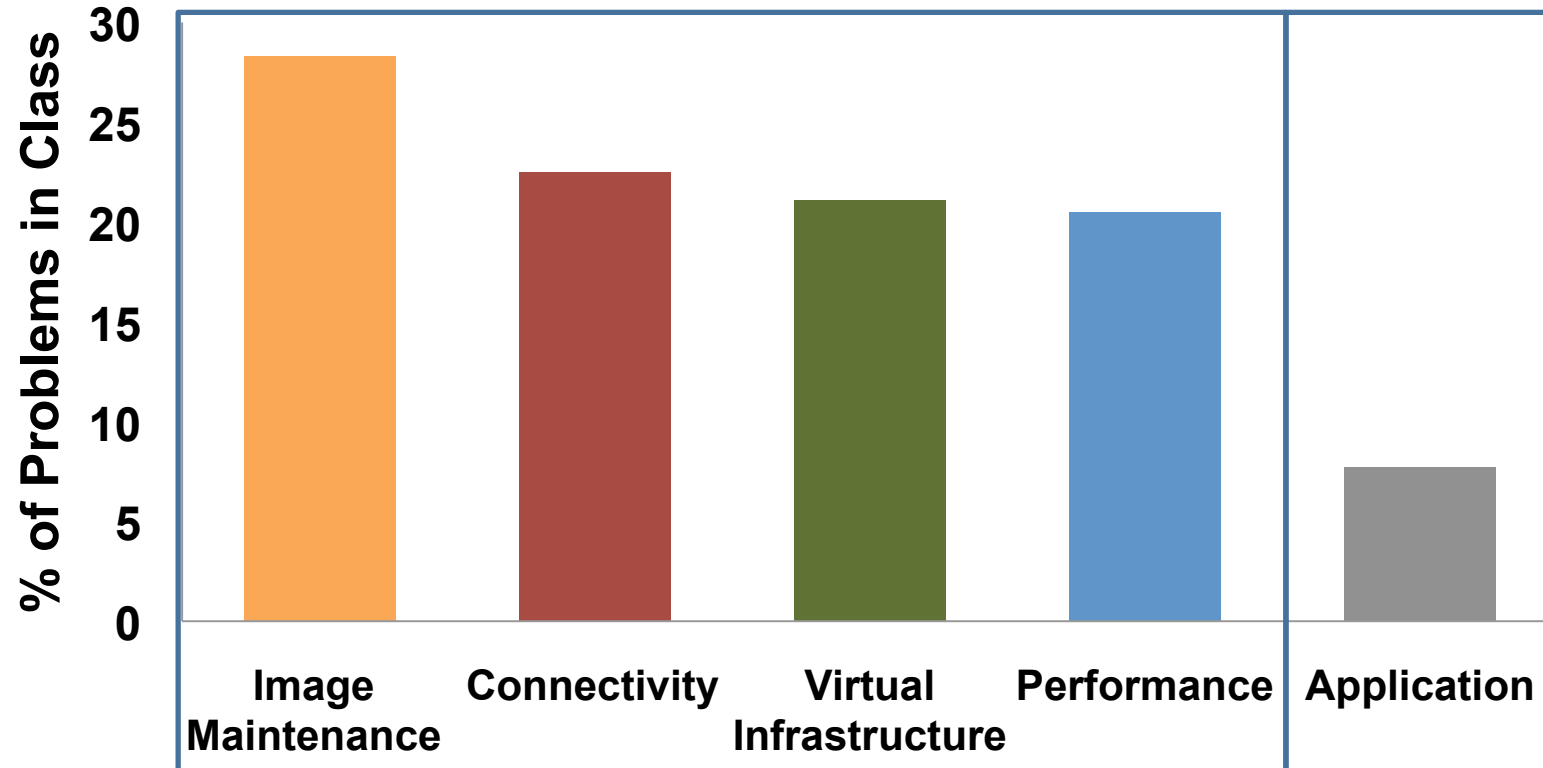
Observed problem taxonomy



Observed problem taxonomy

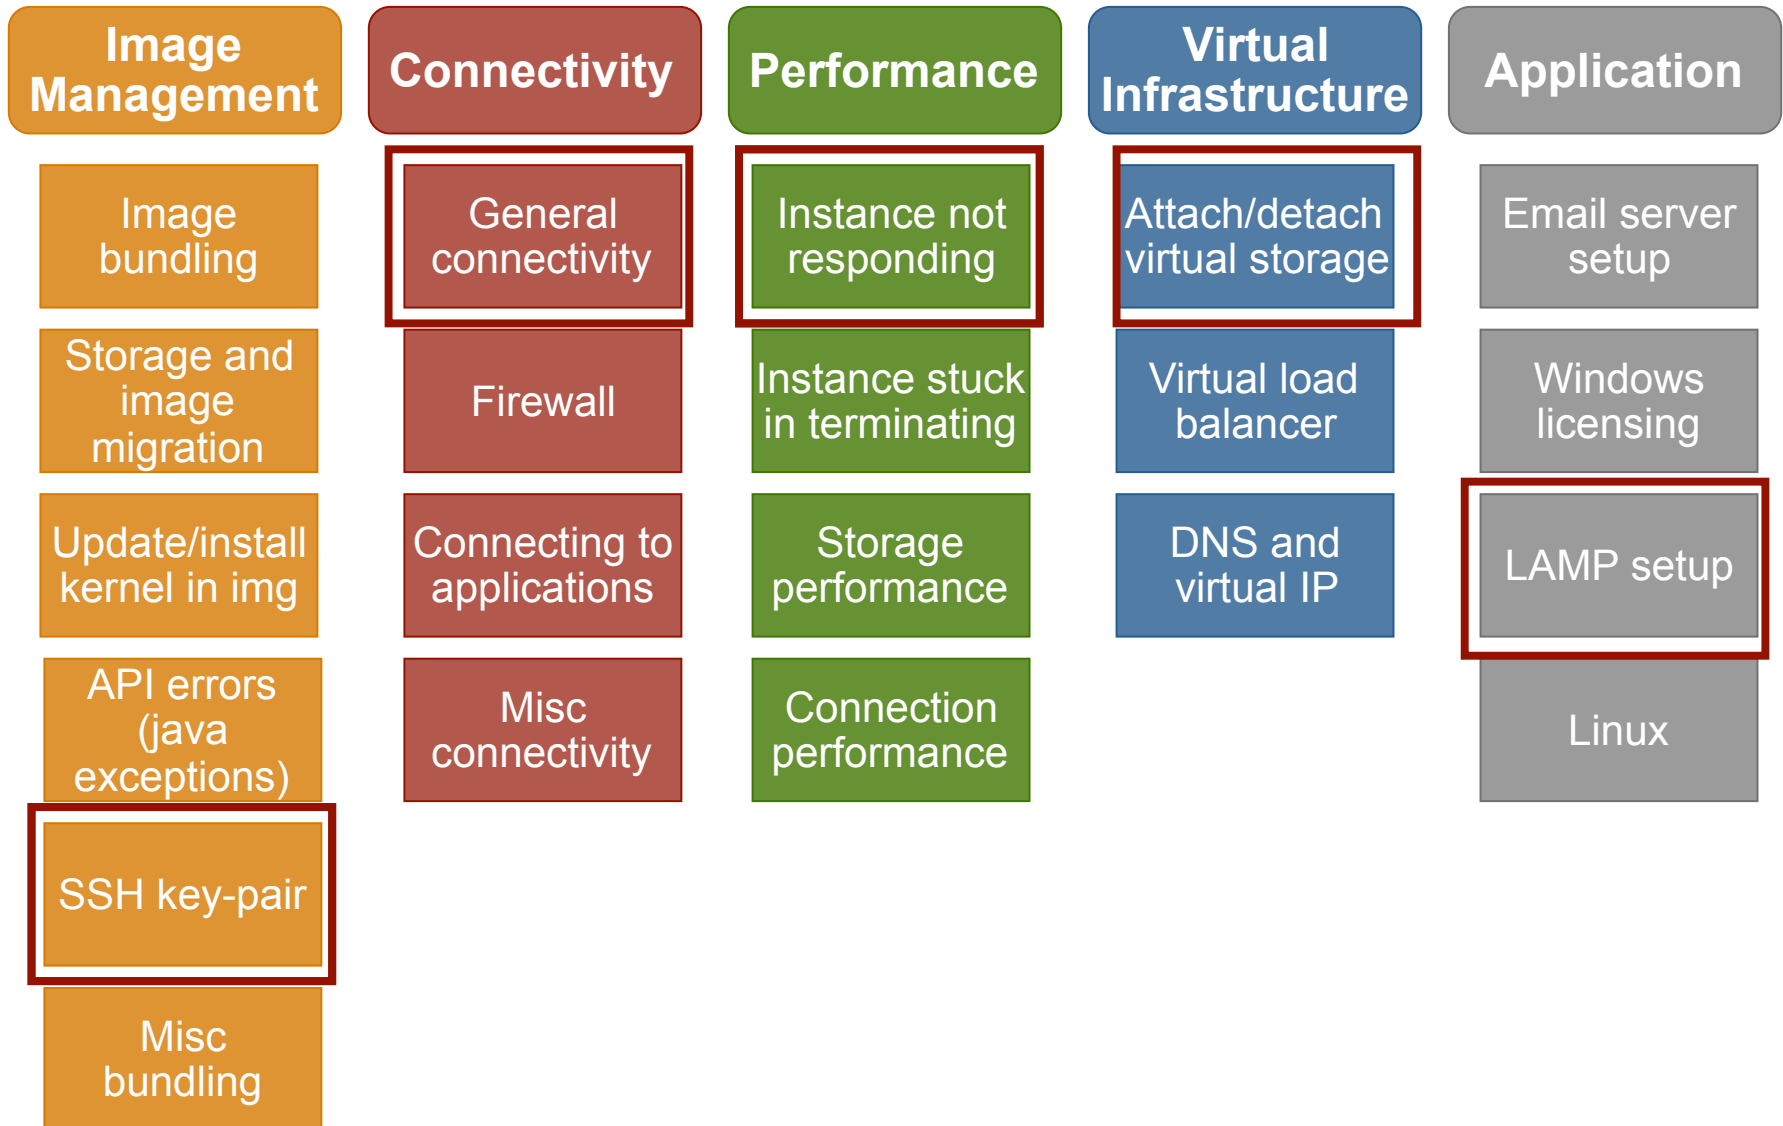


Size of the Different Classes




- Roughly even split across 1st 4 classes
- Fewer application problems fewer (typically out of scope for IaaS)

Observed problem taxonomy

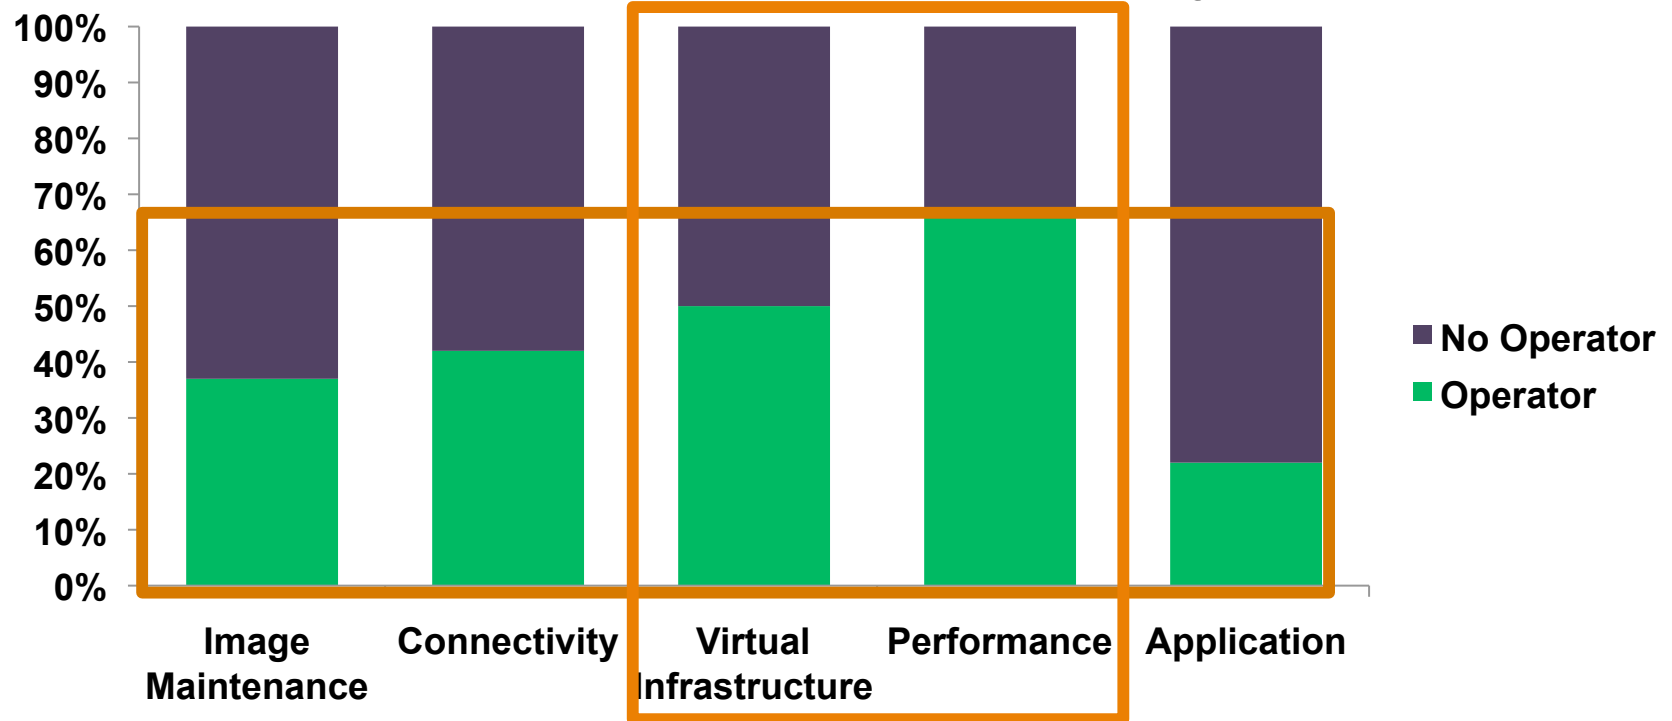


Evolution of Problem Classes




- Image maintenance shrinks: release of new tools
 - Improves on existing techniques
 - E.g 1st Quarter 2007: image manipulation feature
- Virtualization grows: release of new features
 - Users unfamiliar with new feature
 - E.g 3rd Quarter 2008: virtual storage
- Connectivity remains stable

Problem Difficulty



- Users consult operators for difficult problems
- Operators involvement needed for 20-65% of the threads
- Performance & virtual infrastructure requires the most help
 - Performance: virtualization abstracts away details needed to debug
 - Virtual Infrastructure: user can't change state of provider's infrastructure

Evolution of Problem Difficulty



- Over time less help is needed from the operators
 - Forum builds up a DB of solutions
- In a few cases users still need help from the operators
 - Resource abstraction limits user control/visibility
- The virtual infrastructure class is the exception
 - New features are added: users need help understanding them

Observations and Suggestions

- Large fraction of problems solved through self-diagnosis by users
- Limited user visibility causes more operator involvement
 - E.g., Difficult to determine root cause of instance unresponsive
 - Suggestion solution: expose information to the user
- Additional user control for some services could streamline resolution
 - E.g., inspect or change storage volume state
 - Suggestion solution: expose additional control to the user
- New features often introduce new problems
 - Suggestion: Proactive release of companion debugging tools

Some Issues in Enabling Suggested Solutions

- Expose information to the user
 - Protect provider infrastructure details
 - Expose only information relevant to user
 - Reduce collection/storage overhead
- Expose additional user control
 - Multitenancy complications
 - Preventing misuse at the infrastructure level
 - Making the controls scalable

Conclusion

- Empirical study of problems in large provider
 - Classified problem reports into 5 main groups
 - Most persistent problems: virtual infrastructure & performance
 - Problems become less difficult as the forum builds up a database of solutions
- Suggested mechanisms for more effective resolution
 - Expose more information
 - Add user control for infrastructure services
 - Develop specific debug tools for new releases

Thank You

tbenson@cs.wisc.edu

Questions?

Backup

Class 1: Virtualized Infrastructure

- Virtualized components attached to instance
 - Storage, Load balancers
- Sample problem:
 - Unable to detach storage from an instance because of transient hardware failed during detach operation

Class 2: Image Maintenance

- Pertaining to maintenance of an image
 - Bundling/Storing the image
 - Updating or patching image OS
- Sample thread
 - Image throws exceptions during boot up because the image was not built with the appropriate libraries

Class 3: Connectivity

- Connecting to the instance
 - Firewall configuration
 - Connectivity to application
- Sample problem:
 - Unable to connect to instance because kernel ran out of memory and took down the SSH daemon

Class 4: Performance

- General performance of resources in the cloud
 - Performance of instances
 - Performance of virtualization components
- Sample thread
 - User experience poor storage performance because hardware is degraded

Class 5: Application Related

- Related to applications run within the cloud
 - Email server setup/maintenance
 - Setup of LAMP & other issues
 - Windows licensing
- Sample thread
 - How to configure LAMP setup in the presence of dynamic IP addresses