CLAIMS

1. One or more non-transitory computer-readable media storing computer-executable instructions, when executed, cause one or more processors to perform acts comprising:

receiving over a network, from a computing device of an external user of a resource provisioning service comprising multiple, different network resources, a template comprising pre-populated instructions effective at least in part to provision a stack of different types of the multiple, different network resources according to one or more parameters associated with each of the multiple, different network resources, wherein the one or more parameters are specified in the template received from the computing device of the external user;

storing the template at least partly in response to the receiving;

receiving, from the computing device of the external user, a request to provision the stack of multiple, different network resources with use of the template;

retrieving and parsing the template to identify each of the multiple, network resources requested for provisioning and each of the associated one or more parameters; and

at least partly in response to the parsing, attempting to provision each of the multiple, different network resources according to the associated one or more parameters on behalf of the user, wherein the attempt to provision each of the multiple, different network resource is based at least in part on the template received from the computing device of the external user.

2. One or more non-transitory computer-readable media as recited in claim 1, wherein the different types of network resources comprise a storage service, a load balancing service, a compute service, or a security service.

3. One or more non-transitory computer-readable media as recited in claim 1, wherein the multiple, different network resources form at least a portion of a cloud-computing platform.

4. One or more non-transitory computer-readable media as recited in claim 1, wherein each of the one or more parameters associated with a respective one of the multiple, different network resources specifies one of a requested size, a capability, or a characteristic of the respective network resource.

5. One or more non-transitory computer-readable media as recited in claim 1, wherein the template identifies at least one of the multiple, different network resources by a logical name specified by the user.

6. One or more non-transitory computer-readable media as recited in claim 1, wherein the template specifies a particular version of multiple different versions for at least one of the multiple, different network resources.

7. One or more non-transitory computer-readable media as recited in claim 1, wherein the attempting to provision each of the multiple, different network resources comprises attempting to provision at least two of the multiple, different network resources in parallel.

8. One or more non-transitory computer-readable media as recited in claim 1, wherein the attempting to provision each of the multiple, different network resources comprises attempting to provision each of the multiple, different network resources in an order based at least in part on specified dependencies between a portion of the multiple, different network resources.

9. One or more non-transitory computer-readable media as recited in claim 8, wherein the specified dependencies comprise implicit dependencies, and wherein the parsing of the template is effective to identify the implicit dependencies between the portion of the multiple, different network resources.

10. One or more non-transitory computer-readable media as recited in claim 8, wherein the specified dependencies comprise explicit dependencies specified by the user.

11. One or more non-transitory computer-readable media as recited in claim 1, further storing computer-executable instructions, when executed, cause the one or more processors to perform an act comprising notifying the user the stack has been provisioned at least partly in response to successfully provisioning each of the multiple, different network resources.

12. One or more non-transitory computer-readable media as recited in claim 1, further storing computer-executable instructions, when executed, cause the one or more processors to perform acts comprising:

receiving an indication provisioning of one of the multiple, different network resources has failed; and

rolling back each successfully provisioned one of the multiple, different network resources at least partly in response to the receiving of the indication.

13. One or more non-transitory computer-readable media as recited in claim 12, further storing computer-executable instructions, when executed, cause the one or more processors to perform an act comprising notifying the user provisioning of the stack has failed at least at least partly in response to the receiving of the indication.

14. A method, comprising:

under control of one or more computer systems configured with executable instructions,

receiving over a network, from a computing device of an external user of a resource provisioning service comprising multiple, different network resources, a template, wherein the template includes pre-populated instructions for requesting to provision a stack of at least two different types of the multiple, different network resources according to one or more parameters associated with each of the at least two different types of the multiple, different network resources, wherein the one or more parameters are specified in the template received from the computing device of the external user;

parsing the template to identify each of the multiple, different network resources requested for provisioning;

attempting to provision each of the at least two different types of the multiple, different network resources at least partly in response to the parsing, wherein said attempting is based at least in part on the template received from the computing device of the external user;

in response to a case in which each of the at least two different types of the multiple, different network resources were successfully provisioned for the external user in accordance with the template, notifying the external user the stack has been successfully provisioned; and

in response to a case in which one or more of the at least two different types of the multiple, different network resources were not successfully provisioned, notifying the external user the stack has not been successfully provisioned.

15. A method as recited in claim 14, wherein each of the multiple, different network resources reside on a cloud-computing platform.

16. A method as recited in claim 14, further comprising maintaining a user account associated with the user, and wherein the attempting to provision the multiple, different network resources comprises attempting to provision the multiple, different network resources on behalf of the user, wherein a cost associated with a successful provisioning of a network resource is charged to the user account associated with the user.

17. A method as recited in claim 14, wherein the template comprises a text file comprising text-based instructions including calls to each of the multiple, different network resources.

18. A method as recited in claim 14, wherein the template specifies a particular version of multiple different versions for at least one of the multiple, different network resources.

19. A method as recited in claim 14, wherein the template specifies a logical name selected by the user for a component of one of the multiple, different network resources.

20. A method as recited in claim 14, wherein the parsing of the template comprises parsing the template to identify implicit dependencies amongst the multiple, different network resources, and wherein the attempting to provision the multiple, different network resources comprises attempting to provision the multiple, different network resources in an order based at least in part on the implicit dependencies.

21. A method as recited in claim 14, wherein the request explicitly specifies a dependency between at least two of the multiple, different network resources, and wherein the attempting to provision the multiple, different network resources comprises attempting to provision the multiple, different network resources in an order based at least in part on the dependency.

22. A method as recited in claim 14, wherein the attempting to provision the multiple, different network resources comprises attempting to provision at least some of the multiple, different network resources in parallel.

23. A method as recited in claim 14, further comprising tearing down those ones of the multiple, different network resources successfully provisioned when one or more of the multiple, different network resources is not successfully provisioned.

24. An apparatus, comprising:

one or more processors;

memory, accessible by the one or more processors; and

a template-based stack builder stored in the memory and executable on the one or more processors to:

receive over a network, from an external user of a resource provisioning service comprising distinct network resources, a template comprising prepopulated instructions effective at least in part to provision a stack of different types of the distinct network resources based at least in part on a template, the template comprising pre-populated instructions for creating

the stack including the distinct network resources according to one or more parameters specified in the template received from a computing device of the external user;

parse the template to identify each of the distinct network resources; and

attempt to provision each of the distinct network resources on behalf of the external user, wherein the attempt to provision each of the distinct network resource is based at least in part on the template received from the computing device of the external user.

25. An apparatus as recited in claim 24, wherein the template-based stack builder is further executable on the one or more processors to:

notify the user the stack has been successfully provisioned when the template-based stack builder has successfully provisioned each of the distinct network resources on behalf of the user; and

notify the user the stack has not been successfully provisioned when the template-based stack builder fails to successfully provision at least one of the distinct network resources on behalf of the user.

26. An apparatus as recited in claim 25, wherein the template-based stack builder is further executable on the one or more processors to:

roll back each of the distinct network resources the template-based stack builder has successfully provisioned when the template-based stack builder fails to successfully provision the at least one of the distinct network resources on behalf of the user.

27. An apparatus as recited in claim 24, wherein the template-based stack builder is further executable on the one or more processors to receive a request to create the stack for a web application, and wherein the distinct network resources comprise a cloud-based load balancing service, a cloud-based compute service, and a cloud-based storage service.

28. An apparatus as recited in claim 24, wherein the distinct network resources form at least a portion of a cloud-computing platform managed by a single entity.

39. An apparatus as recited in claim 24, wherein one of the distinct network resources forms at least a portion of a cloud-computing platform managed by a first entity and another of the distinct network resources forms at least a portion of a cloud-computing platform managed by a second entity that is independent of the first entity.

30. An apparatus as recited in claim 24, further comprising a database storing or having access to multiple different templates, each of the multiple different templates comprising pre-populated instructions for creating a stack including distinct network resources.

31. An apparatus as recited in claim 31, wherein at least a portion of the multiple different templates have been provided by a community of independent users.

32. A method as recited in claim 14, wherein each of the one or more parameters associated with a respective one the at least two of the multiple, different network resources specifies one of a requested size, a capability, or a characteristic of the respective network resource.

33. An apparatus as recited in claim 24, wherein each of the one or more parameters associated with a respective one the distinct network resources specifies one of a requested size, a capability, or a characteristic of the respective network resource.

34. The method as recited in claim 14, wherein the different types of network resources comprise a storage service, a load balancing service, a compute service, or a security service.

35. The apparatus as recited in claim 24, wherein the distinct types of network resources comprise a storage service, a load balancing service, a compute service, or a security service.

CLAIMS

WHAT IS CLAIMED IS:

1. A method comprising:

under control of one or more computer systems configured with executable instructions,

receiving, from a user, a request to provision a stack of multiple, different network resources identified by a template;

parsing the template to identify each of the multiple, different network resources requested for provisioning;

attempting to provision each of the multiple, different network resources at least partly in response to the parsing;

when each of the multiple, different network resources are successfully provisioned, notifying the user that the stack has been successfully provisioned; and

when one or more of the multiple, different network resources is not successfully provisioned, notifying the user that the stack has not been successfully provisioned.

2. A method as recited in claim 1, further comprising maintaining a user account associated with the user, and wherein the attempting to provision the multiple, different network resources comprises attempting to provision the multiple, different network resources on behalf of the user such that a cost associated with a successful provisioning of a network resource is charged to the user account associated with the user.

3. A method as recited in claim 1, wherein the template includes pre-populated instructions for requesting to provision at least two of the multiple different resources.

4. A method as recited in claim 1, wherein the template comprises a text file comprising text-based instructions that include calls to each of the multiple, different network resources.

5. A method as recited in claim 1, wherein the template specifies a particular version of multiple different versions for at least one of the multiple, different network resources.

6. A method as recited in claim 1, wherein the template specifies a logical name for a component of one of the multiple, different network resources that has been selected by the user.

7. A method as recited in claim 1, wherein the parsing of the template comprises parsing the template to identify implicit dependencies amongst the multiple, different network resources, and wherein the attempting to provision the multiple, different network resources comprises attempting to provision the multiple, different network resources in an order based at least in part on the implicit dependencies.

8. A method as recited in claim 1, wherein the request explicitly specifies a dependency between at least two of the multiple, different network resources, and wherein the attempting to provision the multiple, different

network resources comprises attempting to provision the multiple, different network resources in an order based at least in part on the dependency.

9. A method as recited in claim 1, further comprising tearing down those ones of the multiple, different network resources that have been successfully provisioned when one or more of the multiple, different network resources is not successfully provisioned.

10. An apparatus comprising:

one or more processors;

memory, accessible by the one or more processors; and

a template-based stack builder stored in memory and executable on the one or more processors to:

receive, from a user, a request to create a stack of distinct network resources based at least in part on a template, the template comprising pre-populated instructions for creating the stack that includes the distinct network resources;

parse the template to identify each of the distinct network resources; and

attempt to provision each of the distinct network resources on behalf of the user.

11. An apparatus as recited in claim 10, wherein the template-based stack builder is further executable on the one or more processors to:

notify the user that the stack has been successfully provisioned when the template-based stack builder has successfully provisioned each of the distinct network resources on behalf of the user;

notify the user that the stack has not been successfully provisioned when the template-based stack builder fails to successfully provision at least one of the distinct network resources on behalf of the user.

12. An apparatus as recited in claim 11, wherein the template-based stack builder is further executable on the one or more processors to:

roll back each of the distinct network resources that the template-based stack builder has successfully provisioned when the template-based stack builder fails to successfully provision the at least one of the distinct network resources on behalf of the user.

13. An apparatus as recited in claim 10, wherein the request comprises a request to create the stack for a web application, and the distinct network resources comprise a cloud-based load balancing service, a cloud-based compute service, and a cloud-based storage service.

14. An apparatus as recited in claim 10, wherein one of the distinct network resources forms at least a portion of a cloud-computing platform managed by a first entity and another of the distinct network resources forms at least a portion of a cloud-computing platform managed by a second entity that is independent of the first entity.

15. An apparatus as recited in claim 10, further comprising a database storing or having access to multiple different templates, each of the multiple different templates comprising pre-populated instructions for creating a stack that includes distinct network resources.