

Claims 1-3, 5, 7	Claims 9-11, 13
[1.1] A distributed computing cluster comprising:	[9.1] A distributed computing cluster comprising:
[1.2] a first plurality of computer systems,	[9.2] a first plurality of computer systems,
[1.2.a] wherein each respective computer system of the first plurality of computer systems comprises a memory that stores a respective first data,	[9.3] wherein each respective computer system of the first plurality of computer systems comprises a respective memory that stores a respective first data that: describes at least a respective portion of a configuration of a first transformation pipeline, and identifies at least a second transformation pipeline,
[1.2.b] wherein the respective first data represents a respective portion of a distributed computational graph,	
[1.2.c] and wherein the distributed computational graph describes a flow of output data of a first transformation pipeline to an input of a second transformation pipeline,	
[1.3] wherein a first computer system of the first plurality of computer systems is configured to:	[9.4] wherein a first computer system of the first plurality of computer systems is configured to:
[1.3.a] receive a first stream of input data from a first input feed,	[9.4.a] receive a first stream of input data from a first input feed,
[1.3.b] process the first stream of input data substantially in real time by executing software instructions that apply the first transformation pipeline to the first stream of input data to generate first pipeline output messages,	[9.4.b] process the first stream of input data substantially in real time by executing software instructions that apply the first transformation pipeline to the first stream of input data to generate first pipeline output messages,
[1.3.c] process the respective first data stored in the memory of the first computer system to	[9.4.c] process the respective first data to

Claims 1-3, 5, 7	Claims 9-11, 13
determine information about the second transformation pipeline,	determine an identification of the second transformation pipeline,
[1.3.d] and transmit the first pipeline output messages to a second computer system of the first plurality of computer systems in accordance with the determined information,	[9.4.d] and transmit the first pipeline output messages to a second computer system of the first plurality of computer systems in accordance with the determined identification,
[1.4] wherein the second computer system is configured to:	[9.5] wherein the second computer system is configured to:
[1.4.a] receive the first pipeline output messages,	[9.5.a] receive the first pipeline output messages,
[1.4.b] and process the first pipeline output messages substantially in real time by executing software instructions that apply the second transformation pipeline to the first pipeline output messages to generate second pipeline output messages,	[9.5.b] and process the first pipeline output messages substantially in real time by executing software instructions that apply the second transformation pipeline to the first pipeline output messages,
[1.5] wherein the first and second computer systems are distinct; and	[9.6] wherein the first and second computer systems are distinct; and
[1.6] a second plurality of computer systems;	[9.7] a second plurality of computer systems;
[1.7] wherein a third computer system of the first plurality of computer systems is configured to execute software instructions that cause a fourth computer system of the second plurality of computer systems to execute software instructions that apply at least one of the first transformation pipeline and the second transformation pipeline.	[9.8] wherein a third computer system of the first plurality of computer systems is configured to execute software instructions that cause the second plurality of computer systems to execute at least one of the first transformation pipeline and the second transformation pipeline.
[2] The distributed computing cluster of claim 1, wherein at least one of the first transformation pipeline and the second transformation pipeline is non-linear.	[10] The distributed computing cluster of claim 9, wherein one or more of the first transformation pipeline and the second transformation pipeline is non-linear.

Claims 1-3, 5, 7	Claims 9-11, 13
[3] The distributed computing cluster of claim 2, wherein the third computer system of the first plurality of computer systems is configured to: identify a fifth computer system of the first plurality of computer systems, wherein the fifth computer system is not configured to apply the first transformation pipeline to any stream of input data, and cause the fifth computer system to process a second stream of input data received from a second input feed substantially in real time by executing software instructions that apply the first transformation pipeline to the second stream of input data.	[11] The distributed computing cluster of claim 10, wherein the third computer system of the first plurality of computer systems is configured to: identify a fourth computer system of the first plurality of computer systems, wherein the fourth computer system is not configured to apply the first transformation pipeline to any stream of input data, and cause the fourth computer system to process a second stream of input data received from a second input feed substantially in real time by executing software instructions that apply the first transformation pipeline to the second stream of input data.
[5] The distributed computing cluster of claim 2, wherein the third computer system of the first plurality of computer systems is configured to: identify a fifth computer system of the first plurality of computer systems, wherein the fifth computer system is not configured to apply the second transformation pipeline to any pipeline output message, and cause the fifth computer system to process third pipeline output messages substantially in real time by executing software instructions that apply the second transformation pipeline to the third pipeline output messages.	
[7] The distributed computing cluster of claim 1, wherein the first computer	[13] The distributed computing cluster of claim 9, wherein the first computer

Claims 1-3, 5, 7	Claims 9-11, 13
system is further configured to transmit the first pipeline output messages to the fourth computer system of the second plurality of computer systems, wherein the first computer system is distinct from the fourth computer system, and wherein the second computer system is distinct from the fourth computer system.	system is further configured to transmit the first pipeline output messages to a fourth computer system of the second plurality of computer systems, wherein the first computer system is distinct from the fourth computer system, and wherein the second computer system is distinct from the fourth computer system.