



Runtime Governance Body of Knowledge

for

Artificial Intelligence
and Other Autonomous Systems

John M. Willis

ORCID: 0009-0007-5889-9628

AGCP.ai

An initiative of Sustainable Future Tech, Inc.

Version 1.0 — Zenodo Release

July 2026

Contents

1	Purpose, Scope, and Status	12
1.1	Purpose	12
1.2	Purpose of Runtime Governance	12
1.3	Scope	12
1.4	Out of Scope	14
1.5	Relationship to AGCP	15
1.6	Relationship to Runtime Governance Architecture and PBSAI	16
1.7	Relationship to Existing Frameworks and Standards	16
1.8	Body of Knowledge Status	17
1.9	Foundational Principle	18
2	Audience and Professional Role Profiles	19
2.1	Intended Audience	19
2.2	Audience Backgrounds	19
2.3	Professional Role Profiles	20
2.4	Runtime Governance Executive or Sponsor	20
2.5	Runtime Governance Practitioner	21
2.6	Runtime Governance Assessor	22
2.7	Runtime Governance Architect	23
2.8	Runtime Governance Implementer or Operator	23
2.9	AGCP Conformance Assessor	24
2.10	Cross-Role Knowledge Foundation	25
2.11	Professional Backgrounds and Knowledge Development	26
2.12	Professional Competence Objective	28
2.13	Audience and Role Summary	28
3	Professional Competency Model	29
3.1	Purpose of the Competency Model	29
3.2	Competency Progression	29
3.3	Understanding Runtime Governance	30
3.4	Applying Runtime Governance	30
3.5	Assessing Runtime Governance	31
3.6	Designing Runtime Governance	32
3.7	Evaluating AGCP Conformance	33
3.8	Relationship Among Competency Stages	34
3.9	Relationship to the Training Program	34
3.10	Competency Boundaries	35
3.11	Competency Model Summary	35
4	Core Runtime Governance Thesis	36
4.1	Core Thesis	36
4.2	From Model Governance to Action Governance	36
4.3	Governance at the Point of Consequence	37
4.4	Runtime Governance as Execution-Layer Control	37
4.5	Proposed Transitions as Governance Objects	38
4.6	Admissibility as a Runtime Property	38
4.7	Authority Must Be Current at Commitment	39
4.8	Canonical State as Governance Reality	39

4.9	Evidence as a Runtime Governance Requirement	40
4.10	Structural Refusal	40
4.11	Governance Evidence and Replayability	40
4.12	Deterministic Governance	41
4.13	Runtime Governance and Secure-by-Design AI	42
4.14	The Runtime Governance Gap	42
4.15	Runtime Governance as a Control-Assessment Discipline	43
4.16	Core Runtime Governance Thesis Summary	44
5	AGCP Conceptual Vocabulary	44
5.1	Purpose of the Vocabulary	44
5.2	Action	44
5.3	Proposed Transition	45
5.4	Execution	45
5.5	State Transition	45
5.6	Canonical State	46
5.7	Runtime Context	46
5.8	Governance Context Envelope	47
5.9	Governance Compilation	48
5.10	Policy	48
5.11	Constraint	49
5.12	Invariant	49
5.13	Admissibility	49
5.14	Admissible Set	50
5.15	Authorization	50
5.16	Authority	51
5.17	Authority at Commitment	51
5.18	Binding	51
5.19	Commit	52
5.20	Commit Boundary	52
5.21	Commit-Bound Authorization	52
5.22	Execution-Bound Authorization	53
5.23	Policy Decision Point	53
5.24	Policy Enforcement Point	53
5.25	Evaluation	53
5.26	Enforcement	54
5.27	Structural Refusal	54
5.28	Escalation	55
5.29	Governance Receipt	55
5.30	Refusal Record	55
5.31	Lifecycle State	56
5.32	Append-Only Ledger	56
5.33	Derived Lifecycle State	57
5.34	Replayability	57
5.35	Governance Continuity	57
5.36	Governance Lineage	57
5.37	Non-Bypassability	58
5.38	Time-of-Check/Time-of-Use Consistency	58
5.39	Runtime Assurance	58

5.40	Conceptual Vocabulary Summary	59
6	Runtime Governance Architecture	59
6.1	Purpose of Runtime Governance Architecture	59
6.2	Architectural Premise	60
6.3	Architectural Objectives	60
6.4	Architectural Layers	61
6.5	Intent and Policy Layer	61
6.6	Governance Compilation Layer	62
6.7	Proposal and Context Layer	62
6.8	Runtime Governance Control Plane	63
6.9	Commit and Enforcement Layer	63
6.10	Operational Execution Layer	64
6.11	Evidence, Ledger, and Assurance Layer	64
6.12	Five-Stage Governance Pipeline	65
6.13	Stage 1: Identity and Context Formation	65
6.14	Stage 2: Canonical Representation	65
6.15	Stage 3: Rule Evaluation	66
6.16	Stage 4: Commit Semantics	66
6.17	Stage 5: Audit and Ledger Recording	67
6.18	Control Plane and Execution Plane Separation	67
6.19	Runtime Governance and Policy Engines	68
6.20	Runtime Governance and Identity Systems	68
6.21	Runtime Governance and Workflow Systems	68
6.22	Runtime Governance and Audit Logging	69
6.23	Centralized Runtime Governance Architecture	69
6.24	Distributed Runtime Governance Architecture	70
6.25	Federated Runtime Governance Architecture	70
6.26	Multi-Agent Runtime Governance Architecture	70
6.27	Runtime Governance Reference Architecture Components	71
6.28	Runtime Governance Architecture Failure Modes	71
6.29	Architectural Assessment Questions	72
6.30	Runtime Governance Architecture Summary	72
7	Governance Compilation and Policy-to-Execution Translation	73
7.1	Purpose of Governance Compilation	73
7.2	Governance Intent	74
7.3	The Policy-to-Execution Gap	74
7.4	Governance Compilation Layer	75
7.5	Compiled Governance Artifacts	76
7.6	Governance Compilation Is Not Enforcement	76
7.7	Governance Compilation Is Not a System Invariant	77
7.8	Governance Source Interpretation	77
7.9	Action-Space Definition	78
7.10	Canonical Action Model	79
7.11	Constraint Translation	79
7.12	Evidence-Requirement Translation	80
7.13	Authority-Derivation Rules	81
7.14	Invariant Definition	81

7.15	Refusal-Rule Translation	82
7.16	Escalation-Rule Translation	82
7.17	Commit-Rule Translation	83
7.18	Receipt and Record Requirements	83
7.19	Traceability	84
7.20	Versioning and Change Control	85
7.21	Domain-Specific Governance Compilation	86
7.22	Tenant-Specific Governance Compilation	86
7.23	Human Governance Roles in Compilation	87
7.24	Governance Compilation Quality Criteria	87
7.25	Governance Compilation Assessment Questions	88
7.26	Governance Compilation Summary	88
8	Proposal, Context, Evidence, and Canonical State	89
8.1	Purpose of Proposal, Context, Evidence, and Canonical State	89
8.2	Governed Action Proposal	90
8.3	Proposal as a Governance Object	90
8.4	Proposal Identity	91
8.5	Proposal Structure	91
8.6	Proposal Completeness	92
8.7	Proposal Validity	92
8.8	Governance Context	93
8.9	Governance Context Envelope	93
8.10	Context and Governance Continuity	94
8.11	Context Is Not Canonical State	95
8.12	Evidence	95
8.13	Evidence References	96
8.14	Qualified Runtime Evidence	96
8.15	Evidence Sufficiency	97
8.16	Evidence Freshness	98
8.17	Evidence Provenance	98
8.18	Evidence Integrity	99
8.19	Evidence Continuity	99
8.20	Canonical State	100
8.21	Canonical State Versus Application State	100
8.22	Canonical State Resolution	101
8.23	Canonical State and Lifecycle State	101
8.24	Canonical State and Commit Evaluation	102
8.25	State Freshness	102
8.26	State Integrity	102
8.27	State Drift	103
8.28	Proposal, Evidence, Authority, and State Binding	103
8.29	Proposal Lifecycle Management	104
8.30	Continuous Governance of Proposals	104
8.31	Proposal Degradation	105
8.32	Abandoned Proposals	105
8.33	Assessment Questions	106
8.34	Proposal, Context, Evidence, and Canonical State Summary	106

9	Admissibility, Authorization, Binding, and Commit	107
9.1	Purpose of Admissibility, Authorization, Binding, and Commit	107
9.2	Admissibility	107
9.3	Admissibility Is Not Execution	107
9.4	Admissibility Inputs	108
9.5	Admissibility Function	108
9.6	Deterministic Admissibility	109
9.7	Current Governance Reality	109
9.8	Authorization	109
9.9	Authorization Is Not Authority at Commitment	110
9.10	Authority	110
9.11	Authority Derivation	111
9.12	Execution-Bound Authorization	112
9.13	Commit-Bound Authorization	112
9.14	Binding	113
9.15	Binding Validation	113
9.16	Admissible Set and Adjudication	113
9.17	Commit	114
9.18	Commit Boundary	114
9.19	Commit as Governance Enforcement	114
9.20	Structural Refusal at Commitment	115
9.21	Escalation and Deferred Adjudication	115
9.22	Lifecycle Eligibility	116
9.23	Time-Bound Authorization	116
9.24	Target Binding	117
9.25	Tenant Binding	117
9.26	Revocation and Authority Termination	117
9.27	Time-of-Check/Time-of-Use Control	118
9.28	Decision Outcomes	118
9.29	Authorization Artifacts	118
9.30	Governance Receipts and Refusal Records	119
9.31	Assessment Questions	120
9.32	Admissibility, Authorization, Binding, and Commit Summary	120
10	Determinism, Replayability, and Safety Properties	120
10.1	Purpose of Determinism, Replayability, and Safety Properties	120
10.2	Deterministic Governance	121
10.3	Determinism and Probabilistic Systems	122
10.4	Deterministic Evaluation Consistency	122
10.5	Canonical Representation	123
10.6	Canonical-State Consistency	123
10.7	Replayability	124
10.8	Replayability and Append-Only History	124
10.9	Derived Lifecycle State	125
10.10	Invariant Preservation	125
10.11	Hard Invariant Failure	126
10.12	Non-Bypassability	126
10.13	Enforcement Integrity	127
10.14	Execution-Target Binding	127

10.15	Time-of-Check/Time-of-Use Control	128
10.16	Commit-Time Revalidation	128
10.17	Terminal Execution	129
10.18	Idempotency and Replay Detection	129
10.19	Tenant Isolation	129
10.20	Provenance Preservation	130
10.21	Governance Attribution	130
10.22	Structural Refusal Integrity	131
10.23	Auditability	131
10.24	Forensic Reconstruction	132
10.25	Governance Stability	132
10.26	Replay by Distribution	132
10.27	Safety-Property Assessment Questions	133
10.28	Determinism, Replayability, and Safety Properties Summary	133
11	Runtime Governance Control Categories	134
11.1	Purpose of Runtime Governance Control Categories	134
11.2	Control Category Overview	135
11.3	Identity Controls	135
11.4	Agent and Non-Human Identity Controls	136
11.5	Tenant Identity and Isolation Controls	136
11.6	Authorization Controls	137
11.7	Delegation Controls	137
11.8	Approval and Quorum Controls	138
11.9	State Controls	139
11.10	Canonical-State Controls	139
11.11	Lifecycle-State Controls	140
11.12	Evidence Controls	140
11.13	Evidence Continuity Controls	141
11.14	Provenance and Attribution Controls	141
11.15	Execution Controls	142
11.16	Commit-Boundary Controls	142
11.17	Structural Refusal Controls	143
11.18	Assurance Controls	144
11.19	Replayability Controls	144
11.20	Governance Receipt and Refusal Controls	145
11.21	Control Interdependencies	145
11.22	Control Failure Examples	146
11.23	Control Assessment Questions	146
11.24	Runtime Governance Control Categories Summary	147
12	Lifecycle Sequence and Governance Flow	147
12.1	Purpose of the Lifecycle Sequence	147
12.2	Lifecycle Phases	147
12.3	Phase 1: Intent and Governance Definition	148
12.4	Governance Intent Definition	149
12.5	Authority Model Definition	149
12.6	Evidence Model Definition	150
12.7	Phase 2: Proposal and Authority	150

12.8	Proposal Formation	151
12.9	Proposal Registration	151
12.10	Initial Lifecycle State	152
12.11	Initial Authority Evaluation	152
12.12	Human-in-the-Loop and Quorum Paths	153
12.13	Phase 3: Continuous Governance	153
12.14	Proposal Degradation	154
12.15	Supersession	154
12.16	Expiration	155
12.17	Evidence Preservation During Continuous Governance	155
12.18	Phase 4: Execution Control and Commit Boundary	156
12.19	Pre-Commit Re-Evaluation	156
12.20	Binding Validation	157
12.21	Commit Authorization	157
12.22	Structural Refusal at Commitment	157
12.23	Commit Execution	158
12.24	Phase 5: Post-Commit and Non-Commit Paths	158
12.25	Governance Receipt Generation	159
12.26	Refusal Record Generation	159
12.27	Execution Receipt Generation	160
12.28	Ledger Recording	160
12.29	Lifecycle-State Derivation	161
12.30	Abandoned Proposal Analysis	161
12.31	Lifecycle Flow Summary	162
12.32	Lifecycle Assessment Questions	162
12.33	Lifecycle Sequence Summary	163
13	Multi-Agent, Distributed, and Cross-System Governance	163
13.1	Purpose of Multi-Agent, Distributed, and Cross-System Governance	163
13.2	Multi-Agent Runtime Governance	164
13.3	Agent Roles in Governed Workflows	164
13.4	Agent Identity	165
13.5	Agent-Bound Authority	165
13.6	Delegated Execution	166
13.7	Governance Lineage Across Agents	166
13.8	Governance Context Portability	167
13.9	Cross-System Governance	168
13.10	Cross-System Action Normalization	168
13.11	Cross-System State Consistency	169
13.12	Distributed Governance Domains	170
13.13	Federated Runtime Governance	170
13.14	Governance Interoperability	171
13.15	Cross-Domain Evidence Recognition	171
13.16	Cross-Domain Authority Recognition	172
13.17	Distributed Ordering	173
13.18	Concurrency and Conflict Management	173
13.19	Cross-Agent Tool Use	174
13.20	Agent-to-Agent Delegation	174
13.21	Multi-Agent Escalation	175

13.22	Governance Receipts Across Systems	175
13.23	Refusal Propagation	176
13.24	Governance Continuity Failures	176
13.25	Assessment Questions	177
13.26	Multi-Agent, Distributed, and Cross-System Governance Summary	177
14	Assessment, Testing, and Conformance Foundations	178
14.1	Purpose of Assessment, Testing, and Conformance Foundations	178
14.2	Assessment, Testing, and Conformance	178
14.3	Runtime-Governance Assessment	179
14.4	Assessment Scope	179
14.5	Assessment Criteria	180
14.6	Evidence-Based Assessment	181
14.7	Evidence Sufficiency	181
14.8	Design Adequacy	182
14.9	Operating Effectiveness	183
14.10	Positive Testing	183
14.11	Negative Testing	183
14.12	Refusal-Path Testing	184
14.13	Commit-Boundary Testing	184
14.14	Bypass Testing	185
14.15	Replay Testing	185
14.16	Conformance Testing	186
14.17	Conformance Levels	186
14.18	Conformance Evidence Package	187
14.19	Test-Case Structure	188
14.20	Requirement-to-Test Traceability	188
14.21	Findings	188
14.22	Assessment Conclusions	189
14.23	Unsupported Claims	189
14.24	Scope Limitations	190
14.25	Objectivity and Independence Considerations	190
14.26	Assessment Competence	191
14.27	Representative Assessment Work Products	191
14.28	Assessment Quality Criteria	192
14.29	Assessment and Conformance Questions	192
14.30	Assessment, Testing, and Conformance Summary	193
15	Standards and Framework Alignment	193
15.1	Purpose of Standards and Framework Alignment	193
15.2	Alignment Principle	194
15.3	Framework Alignment Scope	194
15.4	NIST AI Risk Management Framework	195
15.5	NIST Cybersecurity Framework	195
15.6	NIST SP 800-53	196
15.7	NIST SP 800-53 Overlay Use	197
15.8	ISO/IEC 42001	198
15.9	Secure-by-Design Guidance	198
15.10	MITRE ATLAS	199

15.11	OWASP Agentic AI and LLM Guidance	199
15.12	Cloud and AI Control Frameworks	200
15.13	Identity and Zero Trust Alignment	201
15.14	Audit, GRC, and Assurance Alignment	201
15.15	Secure AI Architecture Alignment	202
15.16	Standards Mapping Method	203
15.17	Direct and Indirect Support	203
15.18	Evidence for Framework Alignment	204
15.19	Limitations of Framework Alignment	204
15.20	Standards Alignment Assessment Questions	205
15.21	Standards and Framework Alignment Summary	205
16	Source Provenance and Traceability	206
16.1	Purpose and Scope	206
16.2	Source Status and Interpretation	206
16.3	Primary Published Sources	207
16.4	PBSAI Sources	210
16.5	Supporting Technical Artifacts	211
16.6	Working and Unpublished Sources	213
16.7	Source Categories and Functions	213
16.8	Body of Knowledge Traceability	214
16.9	Citation and Versioning Principles	215
16.10	Source Provenance and Traceability Summary	216
17	Glossary of Core Terms	216
A	Competency and Professional Use of the Body of Knowledge	239
B	Representative Runtime Governance Work Products	245
C	Suggested Reading	271